Districting that minimizes partisan bias

Aaron Brick ^{1*} & Cameron Brick ^{2,3}

in press 2021, *Humanities and Social Sciences Communications*

¹ Department of Computer Science, City College of San Francisco, 50 Frida Kahlo Way, San Francisco, 94112 California, USA

² Department of Psychology, University of Amsterdam, 1012 WX Amsterdam, Netherlands

³ Department of Psychology, Cambridge University, Downing Street, Cambridge, Cambridgeshire, United Kingdom CB2 3EB

**Corresponding author:* Aaron Brick, abrick@ccsf.edu

Author contributions: AB developed the idea and the math, and CB and AB contributed conceptualizing, writing, and editing.

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Districting that minimises partisan bias

[authors blinded]

April 9, 2021

Abstract

The shapes of electoral districts determine how votes translate into seats. When districts favour certain political parties, electoral results can be disproportionate and the public may lose faith in the political process. Disagreement about appropriate district shapes is subjective, rarely resolved, and often leads to lawsuits. Previously, many authors have called for objective districting criteria. We offer a novel synthesis of models that enables the proactive comparison of district maps, by relating a planar graph partition, the single-member plurality rule, the maximin decision rule, and any agreed measure of partisan bias with a territorial map and historical vote results. Employing historical vote totals avoids the complexity and uncertainty associated with counterfactual models of vote swing. Districting plans could be objectively compared on such criteria as party proportionality or compact shape to reject plans with worse bias. Objective tools to reduce partisan bias in district maps could boost collaborative participation, increase perceptions of fairness and justice, and reduce costs.

1 Introduction

Drawing fair constituency boundaries is an outstanding problem in political science. First-past-the-post elections, in which different maps create different outcomes for the same vote, produce partisan bias. Recent problematic cases include the UK Independence Party, which in 2015 got only 1% as many Parliamentary seats per vote as the Conservatives, and the Pakatan Rakyat of Malaysia, which in 2013 took only 40% of seats despite winning more than half of the total vote. Lopsided partisan results lead to underrepresented and unhappy citizens, and may also make electoral competition less important to public policy, which may weaken democracy.

There is a recent explosion of public and expert interest in district shaping and gerrymandering (Arnold, 2017). A recent, small mathematical workshop on the geometry of redistricting drew a thousand applicants. Increased concern about map drawing has prompted ever-subtler means of creating biased plans. In response, scholars have proposed various approaches for crafting fairer plans, often through automation (Altman, 1997; Bozkaya et al., 2003; Browdy, 1990; Brunell, 2010; Forrest, 1964; Nagel, 1972; Puppe and Tasnádi, 2008). Other areas of districting literature concern themselves with organisational processes instead of the resulting maps (Cain, 2011; Cox, 2006), and with the retrospective detection of unacceptable gerrymanders (Grofman and King, 2007; Stephanopoulos and McGhee, 2015). High courts must regularly consider not only the fairness of electoral maps but also whether such matters can even be suitably resolved in court. Even randomly generated maps create partisan bias (Chen et al., 2013).

Unfortunately, the current methods of reducing electoral bias are inadequate, because district criteria and process interventions do not address the partisan

biases of maps. Therefore, it is critical to establish clear criteria and measures for electoral fairness. Below, we present a novel, formal synthesis of models for the objective identification of fair maps through combinatorial optimisation. We provide a formal model for quantifying partisan effects in hypothetical future elections, and an argument for how citizens could collaborate in its productive use.

2 Fairness in districting

Fairness is often considered the absence of bias. In competitive elections, the key bias is a partisan outcome. Observers are divided over the features of a fair districting plan and over how to operationalise fairness in a districting process. Some have suggested a great variety of tweaks and schemes to improve fairness. Others argue that district-based elections work as designed and are inherently fair.

At the individual level, citizens tend to prefer their current electoral system, perhaps due to a psychological effect called system justification wherein individuals are motivated to defend the status quo (Jost and Hunyady, 2005). This preference is an example of the is-ought fallacy, whereby individuals infer what should be from what is, rather than choosing electoral structures based on social values such as justice.

Debate over district boundaries remains wherever single-member plurality elections are held. The way district vote aggregation affects results can strike people as unfair and prompt reform. Citizens in New Zealand adopted a mixed-member proportional system after suffering two electoral inversions in a row. Elsewhere, numerous regulations and theories have attempted to improve districting maps.

One strand of representative elections is a majoritarian one in which the dominant party receives more seats than they have won votes. Such outcomes may make governments more stable and more accountable (Winer, 2019; Powell Jr and Vanberg, 2000). They achieve this by avoiding the coalition governments, typical under proportional representation, which restrain the options of party leaders. A preference for majoritarianism parallels one for a powerful executive.

In contrast, proportionality is a goal that empowers minority interests. Disproportionality can be considered a key indicator of partisan bias (Schuck, 1990) and may indicate a tyranny of the plurality. Therefore, plurality vote aggregation excludes minority groups (Powell Jr and Vanberg, 2000) and establishes poor electoral accountability. Since party platforms and individual partisan identities generally endure, many voters tend to either win or lose just about every election (Guinier, 1994). The repeat losers are often ethnic minorities (Stephanopoulos, 2013). Sectarian violence in the United States, former Yugoslavia, and Northern Ireland has been blamed on plurality elections and majority dominance (Fisher, 1863; Emerson, 2016).

Proportionality has other advantages besides minimising partisan bias. It provides procedural justice in election outcomes (Dunn, 2012): minority constituencies can form coalitions as needed and so achieve effective representation (Guinier, 1994). Countries that use it have desirable social outcomes (Lijphart, 1999) including gender parity among legislators (Salmond, 2006) and better policy correspondence between voters and representatives (Powell Jr and Vanberg, 2000). Even coalition governments tend to exercise checks and balances to keep elections fairer (Norris et al., 2016) such that one party cannot redesign the voting system to favour itself. Finally, declaring proportionality a goal and achieving it both predict greater trust in the legislature (Dunn, 2012).

We do not resolve the ongoing debate about what degree of disproportionality may be considered fair. However, the popular belief that elections are not fair enough has created a vast body of jurisprudence and literature. Below, we review interventions for bias in order to pick an objective goal for more fair districting.

3 Attempts to reduce bias

3.1 Process interventions

Redistricting practices are often opaque and appear unfair, causing political agitation and legal challenges. A fair process upholds procedural justice, which is associated with institutional legitimacy, and impacts public satisfaction as much as elections' actual outcomes (Lind and Tyler, 1988). Nonpartisan process reforms include litigation, neutral commissions, adversarial games, and automated districting.

In some territories, losing interest groups file legal suits and attempt to demonstrate partisan bias. These suits have produced a notorious regulatory thicket, since judges and juries can never reasonably comprehend all of the factors and compromises which go into drawing a map of districts. Explaining the individual aspects of a map — why lines fall in particular places — requires abstract reasoning, and alleged bias may not be amenable to resolution through litigation at all (Issacharoff, 2002). The entire expensive process might be avoided by prior agreement.

Consistent with Rawls's veil of ignorance, those drawing the map should have no stake in its political outcome. The staff of neutral districting commissions

are supposed to have no incentive to introduce partisan bias. However, there is no proven means of testing this assumption.

The neutral commissions still draw maps with partisan bias (Lowenstein and Steinberg, 1985) and can also be expensive. Neutral commissions were adopted in New Zealand in 1887, but they did not prevent electoral inversions there.

Another option is reliance on a disinterested computer program. Semi-automated districting software is common, but despite the expediency and neutrality of such approaches, there are no territories using fully automated districting. Mexico’s districting process is more automated than other countries’, but its opacity does not exclude subjective decisions in district shaping. The degree of trust in software’s creators is a perennial challenge.

Similar to bipartisan commissions, fair division games (Puppe and Tasnádi, 2008) can produce “bipartisan gerrymanders” by formalising bargaining between two parties. In these games, deeply invested parties express their preferences turn by turn. They can alter their own criteria and strategy at any time. There is no reason to think they will find a suitable compromise or produce an outcome that meets specific criteria.

Agreeing on the goal of districting is therefore more productive than circumscribing the process. The unambiguous definition of a single measurable objective is crucial for an effective plan search. With this guidance, programs can use any techniques at all to assess possible solutions until a time or resource limit is reached. An agreed-upon measure of quality and an objectively best-scoring plan could obviate a districting commission and lawsuits over compliance.

3.2 District result criteria

Traditional districting criteria only regulate the nature of individual districts. In contrast, more useful definitions of partisan bias describe the results of elections in many districts. Criteria such as equal population, communities of interest (e.g., based on urban boundaries), and compactness are widely adopted, but unfortunately do not produce fair elections because they do not measure partisan bias, as explained below.

The idea that all districts should have similar populations is a common assumption. However, votes are ineffective for a political minority in a given district. Districts with equal population retain geographical bias and produce election results that may be seen as unfair (Grofman, 1985).

Another popular idea is that districts ought to give identified social communities their own seats. For example, a community might hold a social identity that binds it more closely to one neighbour than another. However, there is no agreed means of measuring similarity, and all the minority interests dispersed across the map — seniors, advocates of net neutrality, etc. — get no representation.

The goal of district compactness also has enduring support, because it is supposed to prevent districting shenanigans like the original strangely shaped “Gerrymander” or the intricate modern maps that connect faraway cities by narrow strips of roadway. There is a common assumption that round district shapes are unbiased, perhaps due to the aesthetic value of circles. Counterintuitively, compactness is unrelated to geographical bias and increasing it does not increase fairness (Lowenstein and Steinberg, 1985). While noncompact districts could well signal gerrymandering, there is no agreement on which of the many definitions of compactness is most appropriate (Niemi et al., 1990). Counterin-

tuitively, compactness is unrelated to geographical bias and increasing it does not increase fairness (Lowenstein and Steinberg, 1985). Therefore, compactness is not a promising criterion to reduce partisan bias.

Some legal criteria are vague or naïve. Other criteria are covert, such as the protection or sacrifice of particular incumbents. Adding more criteria is generally counterproductive because that greatly increases the computational difficulty of reaching existing goals (Coakley, 2008). Additionally, there is no agreement for how to weight different goals in redistricting. Reform attention has therefore turned away from inventing new criteria and is now focusing on districting processes like the model proposed below.

3.3 Party result criteria

Individual district criteria and districting process interventions are poor proxies for electoral fairness. Targeting a direct measure of fairness across all districts would be more effective. Researchers advocating such a criterion have called it “partisan fairness” (Hirsch, 2009) and “gerrymandering in reverse” (Taylor, 1973). In this section, we describe concepts and measures of bias: partisan symmetry, competitiveness and responsiveness, losing voters, the efficiency gap, and party proportionality.

One leading model for electoral fairness is partisan symmetry (Katz et al., 2020). This principle suggests that two parties should receive the same number of seats if they receive the same proportion of the vote. Each party would experience the same bias, if any, as the competition. The only means to test for symmetry is through counterfactual votes. Many of these forecasts are of extremely implausible events, causing a critical missing data problem. The absence of prece-

dent prevents less likely counterfactual outcomes from being modelled (King and Zeng, 2006); partisan symmetry can only be measured when parties could possibly earn each other’s result (Grofman and King, 2007).

Another indication of partisan symmetry or its absence is given by the “equal weight” or “vote bias” metric of McDonald and Best (McDonald and Best, 2015), which greatly resembles the nonparametric skew of the vote distribution. The authors propose the median and mean votes for a given party as good indicators of both intentional gerrymandering and human geography. Unfortunately, correcting the skew for human geography requires a non-existent collection of neutral districting plans. Further, as this measure increases with disproportionality, it lacks a unique contribution.

Symmetry standards also reinforce two-party dominance by excluding other parties from consideration. To our knowledge no jurisdiction has banned third parties, but denying them representation in districting is functionally similar. The traditional seat-vote curve cannot account for the influence of third parties, and because it is not feasible for minority parties like the Bloc Québécois to win nationwide majorities, whether their outcomes are symmetric with those of a larger party is unknowable. In sum, the criterion of partisan symmetry has distinct disadvantages.

Competitiveness is the idea that minimising lopsided elections is desirable, in part because less competitive districts advantage incumbents. Several U.S. states enacted laws making competitiveness one of their districting criteria. Increasing the number of competitive districts would award a bigger bounty to the party winning the most, therefore increasing majoritarian bias (Buchler, 2007) and unhappy, losing voters (Brunell, 2010). Maximising competitiveness would therefore increase partisan bias. Furthermore, in elections with more than two

parties, it is not clear which parties would be in competition.

Recent attention has focused on the effectiveness gap between parties. A single party’s vote efficiency is the ratio of its effective, seat-winning votes to all its votes (a ratio of one would be proportional). The gap between two parties’ vote efficiencies is a measure of geographical bias (Stephanopoulos and McGhee, 2015) and reflects disproportionality. Efficiency gaps are defined as existing between exactly two parties, so they are only useful in strict duopolies.

Many authors have called for proportionality or near proportionality as a desired outcome (Taylor, 1973; Puppe and Tasnádi, 2015). Several measures of proportionality are available, all independent of the number of parties or seats. Although previous work claimed that proportionality did not meet all desired criteria (Taagepera and Grofman, 2003), Koppel and Diskin demonstrated that a cosine similarity model satisfies even stricter criteria (Koppel and Diskin, 2009). This model from the field of information retrieval compares a vector representing the votes won by each party with one representing their seats, in a space with as many dimensions as there are parties. The cosine of the angle between the seat and vote vectors is one in the case of perfect proportionality. This model is among the best available measures of partisan electoral bias.

In sum, many quantitative outcomes are being discussed, but most have serious flaws or apply only to two-party legislatures. Here, we propose a novel districting model using the example objective of party proportionality; the mathematics can be easily adapted to other quantitative goals. In the next section, we propose a new framework for objectively comparing district maps.

4 Model

We present a novel definition of districting criteria to minimise partisan bias that is suitable for combinatorial optimisation. This system combines the concepts of the planar graph partition, the single-member plurality election model, a measure of partisanship, the maximin decision rule, a territorial map, and historical vote results. Our contribution is the synthesis of previous arguments and the proposal that partisan bias be addressed in a transparent and collaborative manner using an objective model like the one described.

The act of districting is a k -partition of a planar graph, well modelled by edge removal. A specific cut set of removed edges therefore defines k districts on the map. Several well-known heuristics can find weight-balanced, or minimum-cost k -partitions, but since voting behaviour cannot be modelled as node weights we rely on general-purpose metaheuristics. Altman found that the number of possible partitions amounted to the Stirling number of the second kind (Altman, 1997), but that formula enumerates set partitions and so gives an overestimate, mostly counting patchwork, discontinuous plans. Lawson’s branch-and-bound method allows exhaustive enumeration of possible partitions (Lawson, 2010).

Finding the absolute best plan under any criteria is well known to be intractably hard. Any search among districting plans is then an optimisation problem. That search simulates elections in different conditions to measure the performance of proposed plans. Without excluding from consideration other strategies, we adopt the Rawlsian maximin and identify the proposal whose worst performance was least bad.

Partitions are evaluated by their expected political outcomes. Historical vote data offers a distribution of possible outcomes that have already happened. By

plugging historical results into proposed maps, the partisanship of hypothetical electoral outcomes can be assessed. Aggregated votes for current and past candidates reflect local turnout trends and parties' enduring bases of support. In the United States, voting history predicts future voting even better than declared intentions (Rogers and Aida, 2014). We suggest using a corpus of recent historical votes, adjusting as necessary for different election formats. The set of historical votes should be comprehensive within a certain period and should include vote swings. Incorporating too many elections means overfitting to the past and slow adaptation to new circumstances such as changes in voter demographics. Conversely, using too few elections risks overfitting to particular circumstances such as specific candidates. Luckily, vote results are sufficiently consistent to provide stable estimates (Linzer, 2012), and in a stabilising feedback loop, the act of voting itself serves to reinforce voters' partisan identities (Dinas, 2014).

4.1 Planar graph partition

M is the map of the territory to be districted, a connected planar graph. The graph has a set of territorial blocks B as nodes, and their adjacencies as edges A . This model is equivalent to the familiar $G = (V, E)$ form.

$$M = (B, A)$$

Every cut set c of removed edges partitions the graph and defines a set of districts D among the blocks B .

$$\forall c \subset A, A \setminus c \Rightarrow B = \bigcup_{d \in D} d$$

The partition creates districts d , which are disjoint subgraphs.

$$\forall (d, d') \in D \times D : d \neq d', d \cap d' = \emptyset$$

K is the set of known valid k -cut sets c , each of which disconnect exactly k subgraphs.

$$K \subset \{c : |D| = k\}$$

4.2 Plurality district election

V is the set of votes cast in B , the set of all blocks, for P , the set of all parties. An individual vote v supports party v_p in block v_b .

$$\forall v \in V, v_b \in B \wedge v_p \in P$$

V_p is the subset of votes V for party p .

$$\forall p \in P, V_p = \{v \in V : v_p = p\}$$

Votes V_d is the subset of votes V from district d .

$$\forall d \in D, V_d = \{v \in V : v_b \in d\}$$

V_{dp} are the votes for party p cast in district d .

$$\forall (d, p) \in D \times P, V_{dp} = V_p \cap V_d$$

4.3 Simultaneous district elections

The seat in each district d is won by some party p and appears in the party's set of districts won, W_p . Parties win any districts in which they earns a plurality of votes. The case of a tie vote is undefined.

$$\begin{aligned} &\forall d \in D, \exists p^* : d \in W_{p^*} \wedge \\ &(\forall p' \in P : p' \neq p^*, \quad |V_{dp^*}| > |V_{dp'}|) \end{aligned}$$

4.4 Cosine measure of proportionality

s_p and v_p are party p 's seat and vote shares, compositional variables which always add to one. A party is proportionally represented when its seat and vote shares are equal.

$$\forall p \in P, \quad v_p = \frac{|V_p|}{|V|}, \quad s_p = \frac{|W_p|}{|D|}$$

m is the districting objective applied to an election result with vote V in districts D . In this step, many objectives are possible. For an example, we define Koppel and Diskin's cosine measure of disproportionality.

$$m(V, D) = \frac{\sum (v_p s_p)}{\sqrt{\sum v_p^2 \sum s_p^2}}$$

4.5 Maximin decision rule

$\mathcal{H}$ is a representative set of recent historical votes V .

$$\mathcal{H} = \{V_1, \dots, V_n\}$$

l is the least favourable outcome obtained in election simulations in districts D using votes from $\mathcal{H}$.

$$l = \min(\{m(V, D) : V \in \mathcal{H}\})$$

c^* is the cut set defining the districting plan whose least favourable outcome, l_{c^*} , is greatest. The case of a tie between plans is undefined.

$$\exists c^* \in K : (\forall c' \in K, c^* \neq c', l_{c^*} > l_{c'})$$

5 Collaborative districting

In the previous section, we presented a novel means of comparing the inherent bias in proposed district maps. Implementation would be separate from the current paper, whose contribution is argument synthesis, a formal definition of partisan bias, and advice about practical implications below. Public participation in districting has several distinct advantages. First, it can improve the perceived legitimacy of institutions (Altman and McDonald, 2010) and likely improve perceptions of procedural justice, which may decrease discontent and legal challenges. Because anyone can monitor or audit the districting process, the resulting map is more legitimate. In a virtuous cycle, justice fosters cooperation (Blader and Tyler, 2013), and the cooperation of the public fosters the identification of the most just plan. Stakeholders may be less motivated to resist an outcome based on criteria seen to be fair (whether proportionality or other criteria). Therefore, wide participation in map drawing should be encouraged and enabled.

Parallelised work also helps with large search problems. A diligent search on a fixed schedule discovers one fairest plan, and each newly discovered superior plan by definition preempts an older one. The winners’ motivations and alternatives are irrelevant to the judgment of the objective function. Relieving authorities from drawing subjective maps may reduce blame and social injustice (Folger and Cropanzano, 2001).

An agreed objective function such as the proportionality standard could also reduce misunderstanding. By transparently seeking to maximise the agreed function, participants have a safe harbour from accusations of partisanship or malfeasance. Someone working in bad faith can produce poor plans, but cannot

trick others into endorsement because their plans are objectively poor. In sum, it is prohibitively difficult to find a winning plan that also accomplishes some nefarious end.

Despite early attempts, distributed districting has faltered so far. The example of Oregon’s 1960 round of legislative districting is cautionary. Staffers produced a booklet called the “Legislative Reapportionment Do-It-Yourself Kit” for members of the state House and Senate, explaining in detail the legal precedents, maps, and population data to help draw new maps. A districting plan somewhat underrepresenting the large population of Multnomah County was passed and signed by the governor, but the State Supreme Court declared it unconstitutional. The plan authors shared a nuanced understanding of the main goals, but could not determine in advance whether their map would be considered valid. Later examples of public participation demonstrated little institutional transparency and little chance for a submitted map to be adopted (Altman and McDonald, 2014). A new collaborative model could minimise the need for lawsuits or committee hearings.

We suggest that a modern kit consist of well-vetted open source software and data which allows interested users to test and propose plans of known validity. Data files containing the block graph and historical votes are necessary for its operation. Cryptographic signatures establish the reliability of the datasets and the authorship of proposed plans. The kit can generate, refine, validate, and submit plans. A range of modalities accommodates the greatest number of participants at little marginal expense: casual users can use the kit online as a service, and serious users can introduce search techniques of their own choice. Unlike commissions convened anew for each redistricting, the kit must be developed only once, then maintained.

Verifiable design can help maximize quality and public trust. We suggest that the kit be licensed as open source (Altman and McDonald, 2010). Furthermore, its version control system and issue tracker should be open to the public, policies which tend to support trust and participation in its development.

Such a kit can use many search strategies; an agnostic approach was first described by Puppe and Tasnádi (Puppe and Tasnádi, 2015). The random mutation of graph partitions is well known. We suggest the following metaheuristics to guide development: simulated annealing (Browdy, 1990) (quantum computers anneal impressively), tabu search (Bozkaya et al., 2003), genetic algorithms (Chandrasekharam et al., 1993), and parallelised versions of the same.

Participation can be made appealing and rewarding through gamification. Those who find good maps could be motivated by social status and monetary prizes. A nominal fee to deter wasteful plan submissions could fund the bounties. Those searching can either work alone or form pools to share risks and rewards. Public-key cryptography or a blockchain could establish proof of authorship.

6 Limitations

This paper explains why an objective method would be valuable, presents the formulas to calculate partisan bias based on agreed criteria, and discusses the practical implications of implementation and effects on public perceptions such as on procedural justice. The model we propose does not provide a solution for multi-member districts. Demonstrating the method on historical vote data or simulations constitutes an effortful software project and would be valuable in a future paper.

The proposed approach can overfit the district map to past votes that may include the effects of strategic voting. By definition, future votes lie outside the sample: their results may not be as proportional. Monte Carlo sampling of various forecasts could target a broader set of outcomes, at the cost of including a complicated model of vote swing.

This model only reflects those past votes represented in the dataset. One consequence is that new parties are included only slowly; they must consistently win votes before districts can be drawn to accommodate them (Brunell, 2010). A 10-year cycle like that of the United States Census may be too long to represent shifting voter interests.

Districts drawn without traditional criteria are unlikely to follow local boundaries, to be compactly shaped or to be similar in size. As malapportionment creates bias only when it reinforces geographical biases (Grofman et al., 1997), population variation could become more acceptable to voters. However, excessive variation could bias political advertising towards the less populous districts. A population floor could ameliorate this problem, but it is difficult to find consensus on how to balance multiple criteria.

Non-compact districts designed to empower minorities could trigger well-founded criticism that some votes are being deliberately wasted. In order to privilege overall fairness in plurality elections, the interests of many individuals would be sacrificed. This problem already exists, but may be perceived as worse if districts are being explicitly drawn to improve fairness.

7 Conclusion

We proposed and formalized a direct remedy for electoral bias in plurality elections. This method serves a social goal and employs no advanced statistics. It is based on established models but could be easily modified. Given that districts must be drawn, we prefer to minimise bias with the objectively least partisan map. This preferable map can only be found by calculating the bias in many simulated elections, a problem well-suited to distributed computing as opposed to humans disputing proposed shapes. Using agreed criteria to guide map selection in an open, public search could improve both electoral fairness and public trust.

8 Data availability

This project contained no data.

9 Competing interests

The authors declare no competing interests.

10 Acknowledgements

We thank Geoff Allen, Sarah Anderson, Deb Behles, David Brick, Jim Dewar, Colin Kuehl, Courtney Gibbons, Daniel O’Leary, Clément Mouhot, David Perkins, Brian Ampere Smith, Yuri Takhteyev, and two anonymous scholars for

generously reviewing drafts of this work.

References

- Altman, M. (1997). Is automation the answer: The computational complexity of automated redistricting. *Rutgers Computer and Law Technology Journal* 23.
- Altman, M. and M. McDonald (2010). The promise and perils of computers in redistricting. *Duke J. Const. L. & Pub. Pol’y* 5, 69.
- Altman, M. and M. P. McDonald (2014). Public participation GIS: The case of redistricting. In *System Sciences (HICSS), 2014 47th Hawaii International Conference on*, pp. 2063–2072. IEEE.
- Arnold, C. (2017). The mathematicians who want to save democracy. *Nature News* 546(7657), 200.
- Blader, S. L. and T. R. Tyler (2013). How can theories of organizational justice explain the effects of fairness? *Handbook of organizational justice*, 329.
- Bozkaya, B., E. Erkut, and G. Laporte (2003). A tabu search heuristic and adaptive memory procedure for political districting. *European Journal of Operational Research* 144(1), 12–26.
- Browdy, M. H. (1990). Simulated annealing: an improved computer model for political redistricting. *Yale Law & Policy Review* 8(1), 163–179.
- Brunell, T. (2010). *Redistricting and representation: Why competitive elections are bad for America*. Routledge.
- Buchler, J. (2007). The statistical properties of competitive districts: what the central limit theorem can teach us about election reform. *PS: Political Science & Politics* 40(2), 333–337.

- Cain, B. E. (2011). Redistricting commissions: a better political buffer? *Yale LJ* 121, 1808.
- Chandrasekharam, R., S. Subhramanian, and S. Chaudhury (1993). Genetic algorithm for node partitioning problem and applications in vlsi design. *IEEE Proceedings E (Computers and Digital Techniques)* 140(5), 255–260.
- Chen, J., J. Rodden, et al. (2013). Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science* 8(3), 239–269.
- Coakley, J. (2008). Electoral redistricting in Ireland. In *Redistricting in Comparative Perspective*. Oxford University Press.
- Cox, A. B. (2006). Designing redistricting institutions. *Election Law Journal* 5(4), 412–424.
- Dinas, E. (2014). Does choice bring loyalty? Electoral participation and the development of party identification. *American Journal of Political Science* 58(2), 449–465.
- Dunn, K. (2012). Voice and trust in parliamentary representation. *Electoral Studies* 31(2), 393–405.
- Emerson, P. (2016). International peace-making — creating post-conflict structures of government. *Open Journal of Political Science* 6(02), 114.
- Fisher, F. (1863). *The Degradation of Our Representative System and Its Reform*. C. Sherman, Son & Co.
- Folger, R. and R. Cropanzano (2001). Fairness theory: Justice as accountability. *Advances in organizational justice* 1, 1–55.

- Forrest, E. (1964). Apportionment by computer. *The American Behavioral Scientist (pre-1986)* 8(4), 23.
- Grofman, B. (1985). Criteria for districting: A social science perspective. *UCLA L. Rev.* 33, 77.
- Grofman, B. and G. King (2007). The future of partisan symmetry as a judicial test for partisan gerrymandering after LULAC v. Perry. *Election Law Journal* 6(1), 2–35.
- Grofman, B., W. Koetzle, and T. Brunell (1997). An integrated perspective on the three potential sources of partisan bias: Malapportionment, turnout differences, and the geographic distribution of party vote shares. *Electoral studies* 16(4), 457–470.
- Guinier, L. (1994). *The tyranny of the majority: Fundamental fairness in representative democracy*. Free Press New York.
- Hirsch, S. (2009). A proposal for redistricting reform: A model state constitutional amendment. *American Mathematical Society, Washington, DC*.
- Issacharoff, S. (2002). Gerrymandering and political cartels. *Harvard Law Review* 2002, 593–648.
- Jost, J. T. and O. Hunyady (2005). Antecedents and consequences of system-justifying ideologies. *Current directions in psychological science* 14(5), 260–265.
- Katz, J. N., G. King, and E. Rosenblatt (2020). Theoretical foundations and empirical evaluations of partisan fairness in district-based democracies. *American Political Science Review* 114(1), 164–178.
- King, G. and L. Zeng (2006). The dangers of extreme counterfactuals. *Political Analysis* 14(2), 131–159.

- Koppel, M. and A. Diskin (2009). Measuring disproportionality, volatility and malapportionment: axiomatization and solutions. *Social Choice and Welfare* 33(2), 281.
- Lawson, B. (2010). A framework for analyzing redistricting. https://www.researchgate.net/publication/228707602_A_Framework_for_Analyzing_Redistricting.
- Lijphart, A. (1999). *Patterns of Democracy: Government Forms and Performance in Thirty-six Countries*. Yale University Press.
- Lind, E. A. and T. R. Tyler (1988). *The social psychology of procedural justice*. Springer Science & Business Media.
- Linzer, D. A. (2012). The relationship between seats and votes in multiparty systems. *Political Analysis* 20(3), 400–416.
- Lowenstein, D. H. and J. Steinberg (1985). The quest for legislative districting in the public interest: elusive or illusory? *UCLA L. Rev.* 33, 1.
- McDonald, M. D. and R. E. Best (2015). Unfair partisan gerrymanders in politics and law: A diagnostic applied to six cases. *Election Law Journal* 14(4), 312–330.
- Nagel, S. S. (1972). Computers & the law & politics of redistricting. *Polity* 5(1), 77–93.
- Niemi, R. G., B. Grofman, C. Carlucci, and T. Hofeller (1990). Measuring compactness and the role of a compactness standard in a test for partisan and racial gerrymandering. *The Journal of Politics* 52(4), 1155–1181.
- Norris, P., F. Martínez i Coma, A. Nai, and M. Grömping (2016). The year in elections, 2015. *Sydney, University of Sydney*.

- Powell Jr, G. B. and G. S. Vanberg (2000). Election laws, disproportionality and median correspondence: Implications for two visions of democracy. *British Journal of Political Science* 30(3), 383–411.
- Puppe, C. and A. Tasnádi (2008). A computational approach to unbiased districting. *Mathematical and Computer Modelling* 48(9-10), 1455–1460.
- Puppe, C. and A. Tasnádi (2015). Axiomatic districting. *Social Choice and Welfare* 44(1), 31–50.
- Rogers, T. and M. Aida (2014). Vote self-prediction hardly predicts who will vote, and is (misleadingly) unbiased. *American Politics Research* 42(3), 503–528.
- Salmond, R. (2006). Proportional representation and female parliamentarians. *Legislative Studies Quarterly* 31(2), 175–204.
- Schuck, P. (1990). Partisan gerrymandering: A political problem without judicial solution. *Political Gerrymandering and the Courts*, 240–251.
- Stephanopoulos, N. O. (2013). Our electoral exceptionalism. *U. Chi. L. Rev.* 80, 769.
- Stephanopoulos, N. O. and E. M. McGhee (2015). Partisan gerrymandering and the efficiency gap. *U. Chi. L. Rev.* 82, 831.
- Taagepera, R. and B. Grofman (2003). Mapping the indices of seats–votes disproportionality and inter-election volatility. *Party Politics* 9(6), 659–677.
- Taylor, P. J. (1973). Some implications of the spatial organization of elections. *Transactions of the Institute of British Geographers*, 121–136.
- Winer, S. L. (2019). In defense of majoritarianism. revised version of Carleton Economic Papers #17-06.