You Can Have Your Cake and Redistrict It Too

Gerdus Benadè
Boston University, Boston, MA, benade@bu.edu

Ariel D. Procaccia
Harvard University, Cambridge, MA, arielpro@gmail.com

Jamie Tucker-Foltz
Yale University, New Haven, CT, j.tuckerfoltz@yale.edu

The design of algorithms for political redistricting generally takes one of two approaches: optimize an objective such as compactness or, drawing on fair division, construct a protocol whose outcomes guarantee partisan fairness. We aim to have the best of both worlds by optimizing an objective subject to a binary fairness constraint. As fairness constraint we adopt the geometric target, which requires the number of seats won by each party to be at least the average (rounded down) of its outcomes under its worst and best possible partitions of the state.

Like proportionality, geometric targets cannot be guaranteed in the general graph redistricting model. We introduce a simplified model of redistricting, which closely mirrors the classic cake-cutting model. This model has two innovative features. First, in any part of the state there is an underlying "density" of voters with political leanings toward any given party, making it impossible to finely separate voters for different parties into different districts. This captures a realistic geography-related constraint without explicitly needing to specify the detailed geography of the state. Second, parties may disagree on the distribution of voters — whether by genuine disagreement or attempted strategic behavior. In the absence of a "ground truth" distribution, a redistricting algorithm must therefore aim to simultaneously be fair to each party with respect to its own reported data. Our main theoretical result is that, surprisingly, the geometric target is always feasible in this model with respect to arbitrarily diverging data sets on how voters are distributed.

A standard for fairness is only useful if it can be readily satisfied in practice. Empirically, we consider real election data and maps of six US states. Despite the existence of worst-case constructions in which the geometric target can not be guaranteed, in all these practical instances we find that the geometric target is feasible, as our state-cutting model predicts. Moreover, imposing it as a fairness constraint comes at almost no cost to three common optimization objectives.

---

1. Introduction

To be elected to the U.S. House of Representatives, a candidate must win a plurality election in their district. These districts are redrawn every decade based on the most recent census; the composition and creation of districts are governed by both federal and state laws. At the federal level, the Voting Rights Act requires that districts be drawn to allow minority groups to fully participate in the democratic process. Locally, many states expect districts to be contiguous and several require districts to be compact and respect "communities of interest."

These guidelines, however, are often open to interpretation. For example, only six states specify a metric by which compactness is measured; elsewhere the determination of whether or not a district is compact is based on rules of thumb. *Gerrymandering*, named after then-Governor of Massachusetts Elbridge Gerry's 1812 approval of a salamander-shaped district thought to aid his Democratic-Republican Party, is the process of exploiting this flexibility by carefully drawing district boundaries for political gain, for example to protect an incumbent or to benefit (or suppress) a specific class, race or political party. It is widely recognized as a distortion of the democratic system (Allen 2023, Bouie 2023); in recent years, mathematicians and computer scientists have mobilized to help address this issue (Duchin 2018).

One place where scientists can contribute is the design of rigorous methods for drawing electoral district maps, which we refer to as *partitions*. This problem is often approached from an optimization perspective (Garfinkel and Nemhauser 1970, Mehrotra et al. 1998, Shirabe 2005, 2009, Oehrlein and Haunert 2017), which involves setting an objective — such as compactness, or the number of "competitive" districts — and finding the optimal partition satisfying the legal constraints (e.g., contiguity, population equality). However, optimization-based approaches do not necessarily lead to *fair* outcomes that would be acceptable to both major political parties.

Our approach. To address the shortcomings of the pure optimization-based approach, we propose to combine it with ideas from *fair division* (Brams and Taylor 1996, Moulin 2003) in a way that ideally enjoys the best of both worlds. On a high level, we wish to enforce an intuitive yet rigorous notion of fairness that is also binary, in the sense that it either is or is not satisfied — there is no question of degree. One key advantage of such a notion is that it would allow a simple *explanation* of why a partition satisfying it is fair (Procaccia 2019). Among all valid partitions that satisfy the fairness notion, we find one that optimizes a given objective function. This approach — optimizing an objective function subject to a binary fairness guarantee — is akin to recent practical success stories in fair division, such as a rent division algorithm (Gal et al. 2017) that has been used to solve tens of thousands of real-world instances.

A key question, of course, is which fairness notion to use. One natural (albeit flawed) answer is *proportionality*: the number of seats won by each party should be proportional to its statewide support. Unfortunately proportionality is not a feasible standard (Nagle 2017). For example, the Republican party won roughly 32% of the Massachusetts statewide vote in the 2016 presidential election. Proportionality suggests that Republicans should win three (roughly 32%) of the state's nine congressional seats. However, this is impossible: there is no partition of the state into nine districts that complies with Massachusetts' redistricting laws under which the Republican party wins any congressional seats based on this election data (Duchin et al. 2019), as the distribution of Republican-leaning voters across the state is rather homogeneous. This is not necessarily disturbing in and of itself; Supreme Court rulings "clearly foreclose any claim that the Constitution requires proportional representation" (DvB 1986).

Instead, we employ the *geometric target* criterion of Landau and Su (2014). To motivate it from our own viewpoint, imagine a procedure in which a fair coin is flipped, and whichever party wins the coin flip is given absolute power to redistrict the state as they wish (subject to the relevant laws regarding contiguity, population equality *etc.*). This procedure would lead to extremely partisan partitions *ex post*, that is, after the coin is flipped. However, it is certainly impartial *ex ante* (before the coin is flipped), as every party is equally likely to suffer or benefit from it. The geometric target distills the essence of what makes this procedure fair, while avoiding its extreme partisan outcomes: each party must win the expected number of districts it would win under the above procedure, rounded down. In other words, the geometric target is the average, rounded down, of the maximum number of districts the party would win under any partition that satisfies the legal constraints, and the minimum number of districts the party would win under any such partition. Rounding is necessary, since it is impossible to guarantee that two parties each win, say, at least 4.5 districts out of nine. We say that a partition is a *GT partition* if the number of districts each party wins is at least its geometric target.

For example, take the 2011 redistricting of Pennsylvania, which the state's Supreme Court ultimately struck down as unconstitutional and replaced with a remedial plan (LWV 2018). Analysis from FiveThirtyEight (Bycoffe et al. 2018) found that the most pro-Democratic map possible leads to nine Democratic congressional seats (out of 18) while the most pro-Republican map leads to five Democratic seats. Based on this, the geometric target of the Democratic party (the average of their extreme outcomes) is seven districts, compared to the five won under the 2011 plan.

A possible objection is that the guarantee given by the geometric target depends on the underlying election data, which can be another source of contention — what happens if the two parties disagree on which dataset should be used to evaluate targets? One of our conceptual contributions is that we explicitly allow the geometric targets of the two parties to be computed with respect to two different datasets. Thus, no matter whether the discrepancies arise from genuine informational disparities or deliberate attempts to achieve a more desirable outcome by manipulating data, any honest party should be satisfied by the final redistricting outcome. This is analogous to the guarantees of cake-cutting protocols: players may disagree over what parts of the cake are valuable, and the protocol must nevertheless find an allocation that is fair for all players according to their respective valuation functions.

As intuitively appealing as this extension of the geometric target is, however, it would not be useful if it cannot be enforced — and so far there has been scant evidence that it can. Even if it can be enforced, it could conceivably restrict the space of feasible partitions to the point of significantly harming standard optimization objectives like compactness. This motivates our research questions:

Do GT partitions exist in theory and are they feasible in practice? If so, is the geometric target compatible with standard optimization objectives?

The validity of our proposed approach hinges on the answers to both questions being positive.

Our results. Our paper has three key contributions. First, we introduce a novel model of redistricting, which we call the *state-cutting* model. It explicitly frames redistricting as a fair division problem, allowing us to import technical tools and methodologies from fair division and rigorously reason about what fairness criteria are feasible. Our second main contribution is a proof that GT-partitions always exist in this model, supporting its usefulness as a fairness axiom. Finally, our third main contribution is an empirical demonstration that GT-partitions exist in practice as well and come at little cost to other objectives.

In Section 2, we motivate the state-cutting model by first considering an alternative *graph redistricting* model that is very well-studied in the redistricting literature. We show that GT partitions may not exist in the worst-case and provide sufficient conditions under which they do. From these observations, we also extract high-level, informal assumptions about the geographic distribution of voters under which GT partitions should exist. These naturally lead to the state-cutting model, which has these assumptions baked in, allowing us to establish worst-case existence guarantees.

In a bit more detail, in our state-cutting model there are no inherent "geometric" constraints on what districts are allowed; instead, we abstract from real life the key challenge that geometry often presents: that the supporters of the two parties cannot be arbitrarily divided between districts. Thus, in the state-cutting model, every part of the state has an underlying density of support for each party. As the name suggests, to capture these density constraints we draw on ideas from the classic cake-cutting model (Brams and Taylor 1996, Robertson and Webb 1998, Procaccia 2013), where densities are defined on the unit interval $[0, 1]$. Under this interpretation, we conceptualize redistricting as the act of partitioning $[0, 1]$ into *districts*, each of which is a finite union of closed intervals (mirroring the typical assumption about pieces of cake). We emphasize that the topology of the unit interval is not meant in any way to capture the geometry or connectivity constraints present in the real world, merely the density constraints. Accordingly, we do not require connectivity of districts in our state-cutting model. We do, however, return to discuss such issues in the latter sections of our paper, from both theoretical and empirical perspectives.

In Section 3 we formally present our state-cutting model and prove our main theoretical result (Theorem 5), that GT partitions always exist, even when the geometric targets of the two parties are computed with respect to two different pairs of density functions (corresponding to two different datasets). Our result is proved via a novel "cut-and-choose" protocol whereby one party divides a strategically critical subset of the interval into two equal pieces and the other party decides which party controls redistricting over which piece. While the protocol is not completely well-defined with respect to the graph redistricting model, we argue by means of empirical simulation that the key steps could be implemented, suggesting our result has implications for more practical settings beyond the (one-dimensional) state-cutting model.

Indeed, in all of our experiments with real election data, we find that GT partitions always exist. In Section 4 we empirically assess the quality of GT partitions in terms of the optimization objectives of compactness, efficiency gap and the number of competitive districts in six U.S. states. We find that restricting our search to GT partitions rarely leads to a significant decrease in any of the three objectives, regardless of whether or not parties agree on the voter distribution. We conclude that the price of enforcing geometric targets as a notion of fairness is extremely low.

Related work. The connection between redistricting and fair division has inspired several papers that put forward interactive protocols by which the parties take turns splitting the state and choosing pieces (Landau et al. 2009, Landau and Su 2014, Pegden et al. 2017, De Silva et al. 2018, Brams 2020, Tucker-Foltz 2018). Of those, our work is most closely related to that of Landau and Su (2014), who introduced the geometric target. They analyze the *LRY protocol* of Landau et al. (2009), in which a neutral administrator presents both parties with a sequence of bipartitions $(L_1, R_1), (L_2, R_2), \ldots, (L_{m-1}, R_{m-1})$ of the state into two pieces, with each $L_i \subseteq L_{i+1}$. For each bipartition, both parties are asked whether they would rather redistrict L_i into i districts or R_i

into $m - i$ districts, with the other party redistricting the other side. If a point of agreement cannot be found, then there must be a specific i at which both parties would prefer redistricting R_i to L_i, but prefer redistricting L_{i+1} to R_{i+1}, so randomness is used to determine whether to use partition i or $i+1$, and which party controls which piece. Landau and Su observe that, *if* the feasible set of electoral maps is constrained to respect a given bipartition, then *at least one* of the two options the parties are asked to choose between must meet their geometric target. However, this does not imply that the final outcome selected by the LRY protocol satisfies the geometric target itself, even for the party whose preferred choice was selected. Landau and Su acknowledge this shortcoming and informally argue that it is unlikely to cause serious problems in practice, appealing to the random elements of the protocol and the neutrality of the administrator.

De Silva et al. (2018) provide a more rigorous treatment of the theoretical guarantees of the LRY protocol, showing that, in the absence of any geometric constraints, both parties are guaranteed to win at least two seats fewer than their geometric targets. However, under a simple grid-based model with a moderate, plausible compactness constraint, they show that the number of districts won by a party can be arbitrarily far from the geometric target. To the best of our knowledge, our paper presents the first protocol that provably satisfies the geometric targets of both parties under a nontrivial model.

We are not the only work to model partisan support via density functions. In particular, a recent (independent) line of work (Gomberg et al. 2023a, Pancs and Sharma 2025, Gomberg et al. 2023b) posits a state that can be sorted from highest density for one party to highest density for the other. On a technical level, this model is the same as ours except that there is only one common density function (rather than a different one for each party) and this function is assumed to be increasing over the unit interval. These papers focus on issues of individual welfare and representation, whereas we are primarily interested in partisan fairness.

Beyond the fair-division viewpoint, *partisan symmetry* (Grofman and King 2007, Niemi and Deegan 1978, Jackman 1994) and the *efficiency gap* (Stephanopoulos and McGhee 2015) are alternative notions aimed at measuring how partisan a proposed plan is. Partisan symmetry requires that statewide swings toward either party would yield comparable swings in seat shares. These hypothetical swings are typically generated by starting from a real election outcome (or a combination of several) and applying uniform (Butler 1952) or approximately uniform swings (King 1989, Gelman and King 1994) to model changes in voters' political preferences. Practically, uniform swings do not allow for the types of changes in voter preferences that occur in reality, and requiring partisan symmetry under more general models of electoral systems can be infeasible. The efficiency gap measures the net difference in the fraction of each party's wasted votes — every vote cast for the minority in a district is deemed to have been wasted, as are all votes for the majority above the threshold required to win the district. Classic gerrymandering techniques like packing (concentrating a party's supporters in one district) and cracking (splitting a party's supporters into minorities in across many districts) lead to large efficiency gaps. A maximum efficiency gap threshold of 8% has been proposed, although there are instances where this is impossible to attain.

These fairness notions generally have the aim of being fair to *parties*, which is the framework we adopt in this paper as well. On the other hand, some prior works have explicitly considered fairness toward *individuals* (Gomberg et al. 2023a, Agarwal et al. 2022). We focus on fairness toward parties firstly because it is simpler, cleaner, and is already a technically deep problem in our model. Secondly, on a conceptual level we feel that notions of fairness toward parties are more likely to be adopted in a political landscape that is currently dominated by parties.

On the optimization side, there is a long history of using optimization models to construct partitions (Hess et al. 1965, Validi et al. 2022). The aim of our work is not to contribute a method for constructing partitions, rather, it is to establish a fairness constraint which can readily be satisfied within an optimization model. Recent work has studied computational methods for redistricting from the perspective that there is an inherent trade-off between fairness and compactness (Schutzman 2020, Swamy et al. 2023, Gurnee and Shmoys 2021, Borodin et al. 2022). Under cardinal measures of fairness such as proportionality or the efficiency gap, there is a "Pareto-frontier" of optimal partitions, at which improving fairness comes at a cost to compactness, and vice versa. Our approach is fundamentally different because our fairness condition is a binary constraint. Thus, our frontier necessarily has only two points: the most compact partition, and the most compact partition satisfying the geometric targets of both parties. In contrast to the recent work of Schutzman (2020), we find that the trade-off is not significant, which is a testament to the robustness and usefulness of the geometric target as a fairness requirement.

Further afield, the classical cake-cutting problem and its close relatives have received significant attention in computer science in general (Procaccia 2013) and in theoretical computer science in particular (Edmonds and Pruhs 2006, Aziz and Mackenzie 2016, Dehghani et al. 2018, Arunachaleswaran et al. 2019). A strength of our paper is that it provides a fundamentally different view of, and a new application domain for, this well-studied problem.

2. The Graph Redistricting Model

Redistricting is commonly modeled as a graph partitioning problem. There is an underlying graph of indivisible geographic units (e.g., counties, precincts, or census blocks) in which edges represent geographic adjacency. Districts must be connected and have equal population weights. A clean model that has been particularly well-studied in prior literature is the special case when the underlying graph consists of equal-population blocks forming a rectangular grid (DeFord et al. 2021, Cannon et al. 2022, De Silva et al. 2018, Procaccia and Tucker-Foltz 2022, Tapp 2021, Donnay and Kahle 2025) or a subgraph of such a graph (Charikar et al. 2022, Frieze and Pegden 2023).

Formally, we write $N := \{1,2\}$ for the set of parties. (We only concern ourselves with the case of two parties in this paper.) We define an instance of the *graph partitioning problem* to be specified by a simple graph $G = (V(G), E(G))$ (which will be a grid or grid subgraph in all of the constructions in this section), a target number of districts m, a vertex weighting function $\mu : V(G) \to \mathbb{R}_{\geq 0}$, and a set of maps $\{f_i^j \mid i, j \in N\}$. For any subset of vertices $S \subseteq V(G)$, we define $\mu(S) := \sum_{v \in S} \mu(v)$ and assume that the total population weight is normalized to be $\mu(V(G)) = 1$. Each $f_i^j : V(G) \to [0,1]$ gives the fraction of support for party j according to party i at each vertex in G. For $S \subseteq V(G)$ we write

$$v_i^j(S) := \frac{1}{|V(G)|} \sum_{x \in S} f_i^j(x)$$

and assume that

$$\sum_{j\in N} f_i^j(v) = \mu(v)$$

for each $i \in N$. This implies that, for any $S \subseteq V(G)$ and party i,

$$\sum_{j\in N} v_i^j(D) = \mu(S).$$

When $f_1^1 \equiv f_2^1$, we say that parties *agree* on voter distributions.

A *district* is a set of vertices inducing a connected subgraph of G. A *balanced, connected partition* of G is a partition of $V(G)$ into districts of size exactly $\frac{1}{m}$. (For the cleanliness of our theoretical model, we require exactly equal populations, though obviously this must be relaxed slightly in practice, as we do for all of our empirical results.)

To discuss the number of seats won by a party with respect to a partition of $V(G)$ into districts, we are confronted with the technical issue of how to resolve perfect ties. Our solution is to assume that whoever is drawing the electoral districts has the ability to resolve ties in whatever way they wish. In other words, a district partition comes with a built-in tie-breaking rule, so to define a partition, one must not only specify the location of the district, but also who wins each district in the case of a tie. Our results do not depend critically on this modeling choice; it is mainly for elegance and ease of exposition. Formally, for any $m \in \mathbb{Z}_{\geq 1}$ and $S \subseteq V(G)$ whose size $\mu(S)$ is an integer multiple of $\frac{1}{m}$, an *m-partition of S* is a pair (P,T), where $P = \{D_1, D_2, \ldots, D_{m\mu(S)}\}$ is a balanced connected partition and $T : P \to N$ is a tie-breaking rule. We write $\mathcal{P}(m)$ for the set of all m-partitions of $[0,1]$. Given an instance of the graph partitioning problem and an m-partition (P,T), we denote the number of districts won (in the sense of absolute majority) by each party $j \in N$, according to party $i \in N$, by

$$u_i^j(P,T) := \left| \left\{ D \in P \mid v_i^j(D) > \frac{1}{2m} \text{ or } \left(v_i^j(D) = \frac{1}{2m} \text{ and } T(D) = j \right) \right\} \right|.$$

For each district D in the set above, we say that party j *wins D according to i under* (P,T). When $j = i$, we simply say *i wins D under* (P,T). A *GT partition* is an m-partition (P,T) of $[0,1]$ such that, for all $i \in N$, the *geometric target for party i* is satisfied:

$$u_i^i(P,T) \geq \left\lfloor \frac{1}{2} \left(\min_{\substack{(P',T') \\ \in \mathcal{P}(m)}} u_i^i(P',T') + \max_{\substack{(P',T') \\ \in \mathcal{P}(m)}} u_i^i(P',T') \right) \right\rfloor.$$

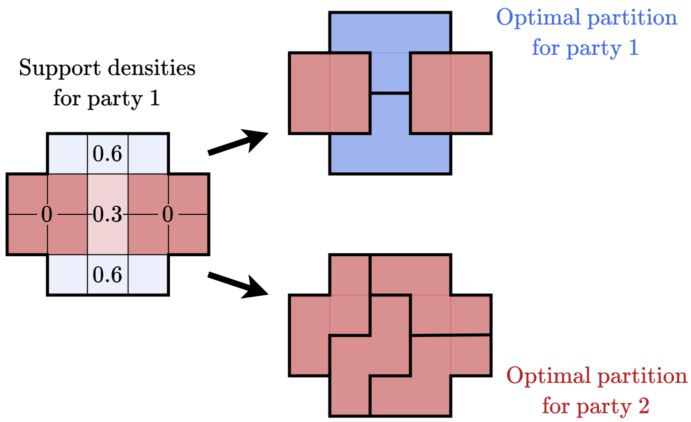

Figure 1 Two instances of the graph partitioning problem on grid subgraphs in which parties agree on voter distributions. The value of $f_1^1 = f_2^1$ (density of support for party 1) is depicted in each region. In the instance on the left, GT-partitions do not exist; in the instance on the right, they do. Note that the partition on the top left is an isolated in the ReCom state-space, also known as a *locked tiling* (Tucker-Foltz 2023), which is why Theorem 2 does not apply.

We seek to find assumptions under which we can guarantee the existence of GT partitions. Unfortunately, we run into the following obstacle:

THEOREM 1. *Even on feasible instances of the graph redistricting problem on grid subgraphs with $m = 4$ districts and $f_1^1 \equiv f_2^1$, a GT partition may not exist.*

Proof. Consider the instance shown on the left of Figure 1. The geometric target is for party 1 to win $\frac{0+2}{2} = 1$ district. One can verify, by enumerating all graph partitions, that there is no partition where party 1 wins exactly 1 district (regardless of the tie-breaking rule, since ties will never occur). □

In the case where parties agree, existence of GT partitions is related to a well-studied question about the *ReCom* Markov chain (DeFord et al. 2021). The state-space of ReCom is the set of all balanced connected partitions a graph. State transitions are called *recombination* moves, in which two random adjacent districts being merged together and split in a different way while preserving the connectivity and balance constraints. When using ReCom to generate partitions on real-world data, as we do in Section 4, it is necessary to set a small tolerance for population imbalance.

However, we presently define ReCom with strict exact balance constraints: every part must have an exactly equal number of vertices. It is widely believed that, for a fixed number of districts m, for sufficiently large grids and well-connected grid subgraphs, ReCom is *ergodic*, meaning that it is possible to get from any balanced connected partition to any other with a finite number of recombination moves.

THEOREM 2. *Given a feasible instance* $(G, m, \{f_i^j\})$ *of the graph redistricting problem, if* $f_1^1 \equiv f_2^1$ *and ReCom is ergodic on* G *with* m *districts, then a GT partition exists.*

Proof. Let (P_1, T_1) be a best m-partition of G for party 1 (which is a worst m-partition for party 2), and let (P_2, T_2) be a worst m-partition of $[0,1]$ for party 1 (which is a best m-partition for party 2). Without loss of generality assume each T_i breaks ties in favor of party i. By ergodicity, there is a chain of partitions from P_1 to P_2, each differing from the previous one on at most 2 districts. Consistently using T_1 to break ties, we get a chain of m-partitions from (P_1, T_1) to (P_2, T_1).

We claim that, at each step in this chain, the number of districts won by party 1 (and thus party 2 as well) changes by at most ± 1. Suppose toward a contradiction that this was not the case at some step, going from (P, T_1) to (P', T_1). Let v^1 denote the common function $v_1^1 \equiv v_2^1$. Let the two districts on which P and P' differ be $D_1, D_2 \in P$ and $D_1', D_2' \in P'$. Since we are breaking ties in favor of party 1, the only way that the number of wins can differ by at least 2 is if party 1 has a weak majority in D_1 and D_2, but a strict minority in D_1' and D_2'; or a strict minority in D_1 and D_2, and a weak majority in D_1' and D_2'. These two cases are completely analogous, so we only consider the former case, i.e., $v^1(D_1) \geq \frac{1}{2m}$, $v^1(D_2) \geq \frac{1}{2m}$, $v^1(D_1') < \frac{1}{2m}$, and $v^1(D_2') < \frac{1}{2m}$. Then, by the additivity of v^1,

$$\frac{1}{m} \leq v^1(D_1) + v^1(D_2) = v^1(D_1 \cup D_2) = v^1(D_1' \cup D_2') = v^1(D_1') + v^1(D_2') < \frac{1}{m}.$$

We have a contradiction, so the number of districts won by party 1 can change by at most ± 1 at each link in the chain.

Finally, we extend the chain by m more steps from (P_2, T_1) to (P_2, T_2) by changing the tie-breaking rule one district at a time. Again, the number of wins for party 1 changes by at most

± 1 at each step. Thus, at some point in the middle of the chain of m-partitions from (P_1, T_1) to (P_2, T_2), the rounded average number of wins for each party between the extremes is realized. □

The main conceptual takeaway from these results is that divisibility of population is crucial. Counterexamples where ReCom is not ergodic arise because there are large geographic areas that cannot be further subdivided. In the counterexample from Theorem 1, for instance, if we subdivide each vertex into 4 equal pieces with the same densities as before, it becomes possible to construct a GT partition as shown on the right of Figure 1. In fact, one of the few positive results about the state-space of ReCom being connected holds in a continuous model where a redistricting map is a partition of a square into polygons of equal area (A. Akitaya et al. 2023). Extrapolating, we may form the following informal hypothesis about the feasibility of the geometric target:

ASSUMPTION 1. *Population is continuously divisible, i.e., we can perturb district boundaries without seeing any large jumps in population.*

HYPOTHESIS 1. *When parties agree on voter distributions, Assumption 1 implies that GT partitions exist.*

Unfortunately, divisibility alone is not sufficient in the general case, when parties do not agree. Even on square grids, it is impossible to obtain any (additive or multiplicative) approximation to the GT utilities for both parties simultaneously.

THEOREM 3. *For any integer $m \geq 2$, there is a feasible instance of the graph redistricting problem on an $m \times m$ grid with m districts on which the geometric target of each party is at least $\lfloor \frac{m}{4} \rfloor$, yet any m-partition gives some party utility zero. This holds even if all vertices are further subdivided into arbitrarily large rectangular grids.*

Proof. For party 1, define the density f_1^1 to be $\frac{1}{2} + \frac{1}{2m}$ on every other row (starting with the first one), and zero on the other rows. For party 2, analogously define the density f_2^2 to be $\frac{1}{2} + \frac{1}{2m}$ on every other column (starting with the first one), and zero on the other columns. Clearly, partitioning the graph into rows gives party 1 utility $\lceil \frac{m}{2} \rceil$ and partitioning the graph into columns gives party 2 utility $\lceil \frac{m}{2} \rceil$. Thus the geometric targets of both parties are at least $\lfloor \frac{m}{4} \rfloor$.

Fix an m-partition (P,T). If any district $D \in P$ includes vertices in multiple rows, then it contains some vertex x for which $f_1^1(x) = 0$. Since f_1^1 is bounded by $\frac{1}{2} + \frac{1}{2m}$, we have

$$v_1^1(D) = \frac{1}{m^2}\left(f_1^1(x) + \sum_{y \in D\setminus\{x\}} f_1^1(y)\right) \leq \frac{1}{m^2}\left(0 + (m-1)\left(\frac{1}{2} + \frac{1}{2m}\right)\right) = \frac{m^2-1}{2m^3} < \frac{1}{2m},$$

so party 1 loses D according to party 1. Thus, for party 1 to have nonzero utility, some district $D \in P$ must be completely contained in the same row, which means it must be an entire row. Symmetrically, for party 2 to have any utility some district $D' \in P$ must be an entire column. But for $m > 2$, D and D' are not the same yet intersect, so P cannot contain them both. Thus, some party must have zero utility.

Finally, in the case where we replace each vertex with a grid of vertices of identical densities, we note that the same argument goes through: It is still not allowable for a district to include any part of another row/column, otherwise the party will lose the district. □

The obstruction here is that the parties wish to form districts in very specific parts of the graph. These regions are not nicely-shaped, which makes it impossible to satisfy both parties' demands simultaneously without breaking connectivity constraints. One might boldly conjecture that this is the *only* remaining obstruction:

Assumption 2. *When looking for a large region where a given party can hope to win many districts, it will be possible to find one that is nicely-shaped.*

Hypothesis 2. *Even when parties disagree on voter distributions, Assumptions 1 and 2 imply that GT partitions exist.*

We argue that this is indeed the case; and moreover, Assumptions 1 and 2 do hold in real world redistricting settings. The main technical tool we use is our *state-cutting* model, introduced in the next section, which comes with these assumptions built in.

3. The State-Cutting Model

We now introduce the state-cutting model for redistricting. We adopt a similar notation as in the graph redistricting model, where $N = \{1, 2\}$ is the set of parties. The state is represented by the interval $[0,1]$. A *district* is a subset of $[0,1]$ that can be expressed as a finite union of closed intervals. An instance of the *state-cutting problem* is specified by a target number of districts $m \in \mathbb{Z}_{\geq 1}$ and a set of *voter distribution functions* $\{f_i^j \mid i,j \in N\}$ giving the density of support for party j according to party i over any district, where each $f_i^j : [0,1] \to [0,1]$. As is customary in the cake-cutting model, for any district D, we write

$$v_i^j(D) = \int_D f_i^j(x)dx.$$

We assume population density has been normalized so that, for any $x \in [0,1]$ and $i \in N$,

$$\sum_{j \in N} f_i^j(x) = 1.$$

This implies that, for any district D and party i,

$$\sum_{j \in N} v_i^j(D) = \mu(D)$$

An m-partition (P,T) of a set $S \subseteq [0,1]$ is defined analogously as before, where $T : P \to N$ is a tie-breaking rule and P is a set of districts such that:

1. For all k, $\mu(D_k) = \frac{1}{m}$.
2. For all k_1, k_2, $\mu(D_{k_1} \cap D_{k_2}) = 0$ (i.e., districts only overlap at endpoints).
3. $\bigcup_k D_k = S$.

We define winning districts, party utilities u_i^j, and geometric targets exactly as in Section 2.

Figure 2 begins a hypothetical running example instance of the state-cutting problem where $m = 10$. In this instance, we may define a 10-partition (P,T) by taking

$$P := \{[0,0.1],[0.1,0.2],[0.2,0.3],[0.3,0.4],[0.4,0.5],[0.5,0.6],[0.6,0.7],[0.7,0.8],[0.8,0.9],[0.9,1]\}.$$

According to party 1, party 1 only wins districts $[0.3,0.4]$, $[0.4,0.5]$, $[0.5,0.6]$, and, depending on T, $[0.2,0.3]$. Party 2 agrees with this assessment, except that party 1 also wins $[0.6,0.7]$ according to party 2. As shown in Section 3.4, the geometric target for party 1 is to win at least $\lfloor \frac{0+8}{2} \rfloor = 4$ districts, and the geometric target for party 2 is to win at least $\lfloor \frac{3+10}{2} \rfloor = 6$ districts, each according to their own voter distribution functions. Thus, if we set $T([0.2,0.3]) := 1$, the geometric target for

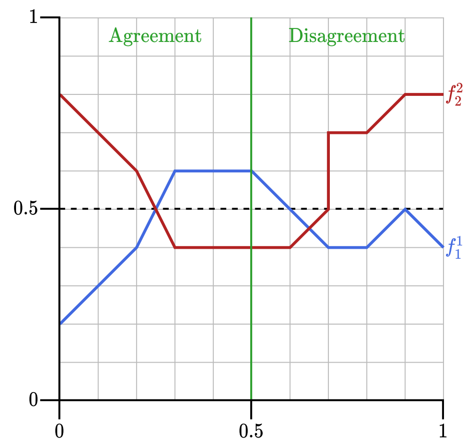

Figure 2 An instance of the state-cutting problem where and $m = 10$. The density functions f_1^1 and f_2^2 are shown in blue and red, respectively. This is a full specification of the instance, since we must have $f_1^2(x) = 1 - f_1^1(x)$ and $f_2^1(x) = 1 - f_2^2(x)$, and the voter distribution functions can be computed by taking integrals, e.g., $v_1^1([0.5, 0.7]) = \int_{0.5}^{0.7} f_1^1(x)dx = 0.1$. The two parties happen to agree on the distribution of voters over $[0, 0.5]$, but disagree everywhere else.

party 1 will be satisfied; if we set $T([0.2, 0.3]) := 2$, the geometric target for party 2 will be satisfied; but there is no choice of tie-breaking rule satisfying both targets simultaneously. In other words, for this choice of P, there is no T such that (P,T) is a GT partition.

Before proceeding, it is worth comparing the state-cutting problem to the traditional cake-cutting problem studied in fair division (Brams and Taylor 1996). The cake is typically modeled as the unit interval, like our state. A *piece of cake* is a finite union of disjoint intervals, which is how we define a district. In the cake cutting model, each agent has a value density function specifying their value for a piece of cake, similar to our voter distribution functions. These similarities will allow us to import the idea of 'cut-and-choose' protocols and the Austin cut procedure (Austin 1982) from the cake-cutting literature. However, there are some crucial differences: in the cake-cutting model the unit interval is divided into as many pieces as there are agents, pieces are assigned to agents and fairness depends on the respective values of the agents for their own pieces. The state-cutting model requires the formation of m districts even for two parties, and it allows each party to report a voter distribution function for every party (not only themselves). Furthermore, districts are not assigned to any party and the value of a district for a party is binary and depends on which party

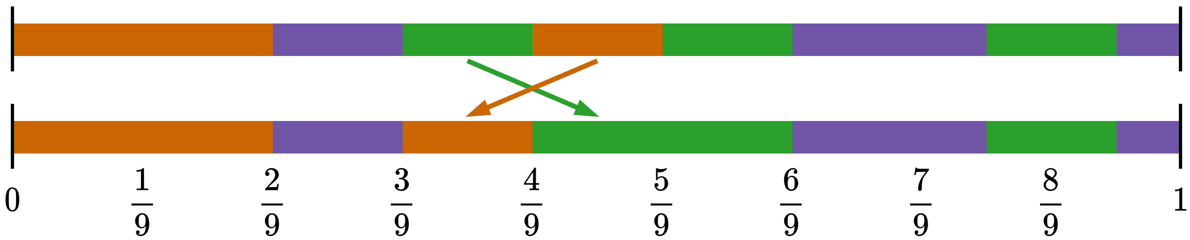

Figure 3 An illustration of the repartitioning step in the proof of Thm. 4 with three districts, where intervals of the same color form a district. The top partition is

$\{[0,\frac{2}{9}]\cup[\frac{4}{9},\frac{5}{9}],[\frac{3}{9},\frac{4}{9}]\cup[\frac{5}{9},\frac{6}{9}]\cup[\frac{15}{18},\frac{17}{18}],[\frac{2}{9},\frac{3}{9}]\cup[\frac{6}{9},\frac{15}{18}]\cup[\frac{17}{18},1]\}$, and the bottom partition is

$\{[0,\frac{2}{9}]\cup[\frac{3}{9},\frac{4}{9}],\quad[\frac{4}{9},\frac{6}{9}]\cup[\frac{15}{18},\frac{17}{18}],[\frac{2}{9},\frac{3}{9}]\cup[\frac{6}{9},\frac{15}{18}]\cup[\frac{17}{18},1]\}$.

has greatest support (i.e. on the voter distribution functions of all parties). Finally, the ‘population balance’ requirement that $\mu(D)=1/m$ rarely has an analogue in cake cutting.

3.1. Existence of GT Partitions: When Parties Agree on Voter Distribution

It is relatively straightforward to see that GT partitions always exist in the case where $f_1^1\equiv f_2^1$. The following theorem is superseded by our main result (Theorem 5 in the next section), but it is nevertheless instructive as a warm-up.

Theorem 4. *Given any instance of the state-cutting problem in which $f_1^1\equiv f_2^1$, a GT partition always exists.*

Proof. We use the same argument as in Theorem 2: It suffices to show that we can connect any pair of partitions P_1 and P_2 of $[0,1]$ into equal-population districts by a sequence of intermediate partitions, each differing from the previous one on at most 2 districts. For any given $i\in\{1,2\}$, we imagine bubble-sorting the disjoint intervals comprising the districts of P_i, where the sort key of an interval is the index of the district in P_i to which it belongs. Each time two adjacent intervals are swapped, we repartition the corresponding subinterval to get a new partition, as in Figure 3. In the end, we arrive at the simplest possible partition P^*, in which each district is connected (like the example P from Section 3). This creates a chain of partitions from P_1 to P^* to P_2, so the same argument from Theorem 2 implies that there is some intermediate m-partition satisfying the geometric targets of both parties. □

Notice that this proof technique does not work in the general case where parties do *not* agree on voter distributions, as the intermediate partitions may yield lower utilities for both parties.

3.2. Existence of GT Partitions: General Case

We now state the main theoretical result of this paper and outline its proof.

THEOREM 5. *In any instance of the state-cutting problem, a GT partition exists.*

The proof is via an interactive protocol, which we call the GT Protocol. It is important to note that we do not suggest using this protocol; rather, it is merely a tool for proving the feasibility of GT partitions, providing theoretical underpinnings for our proposed approach of using geometric target constraints for optimization, as we discuss in Section 4. We motivate the GT Protocol by first considering a simplified version that does *not* work.

Cut-and-Choose Protocol:

1. Party 1 splits $[0,1]$ into two pieces D_1 and D_2, each of measure an integral multiple of $\frac{1}{m}$.
2. Party 2 chooses distinct $k_1, k_2 \in N = \{1,2\}$.
3. Party 1 redistricts D_{k_1}.
4. Party 2 redistricts D_{k_2}. Combine together to get a complete m-partition of $[0,1]$.

To see what goes wrong, consider the instance of the state-cutting problem in Figure 4. The geometric target for party 2 is $\lfloor \frac{10+5}{2} \rfloor = 7$. However, if party 1 divides the state into $D_1 = [0, 0.5]$ and $D_2 = [0.5, 1]$ in Step 1, then both parties will end up with 5 seats regardless of how the rest of the game is played. Hence, party 2 fails to meet its geometric target. This example illustrates the deficiency of the Cut-and-Choose protocol: Unlike the eponymous protocol from cake-cutting, in the state-cutting model the mere act of cutting confers power to the cutter, beyond the subsequent utility the parties gain from redistricting each piece. The GT Protocol, presented below, addresses this issue by only letting the cutter subdivide a smaller region of the state which they consider to be the most important. Whoever desires the region of smaller measure gets to be the cutter. This balances out the relative power of each role.

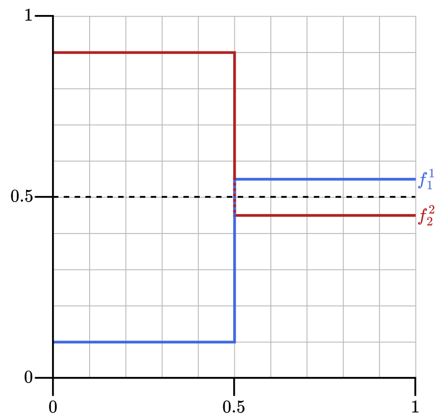

Figure 4 An instance of the state-cutting problem where $m = 10$ in which the simple Divide-and-Choose protocol fails to guarantee party 2 its geometric target. The left half of the state has a strong advantage for party 2 and the right half of the state has a weak advantage for party 1.

GT Protocol:

1. Each party simultaneously draws a district of arbitrary size. Call these districts X_1 and X_2.
2. Let party j (the "cutter") be the party with the smaller district, and party i (the "chooser") be the party with the larger district (i.e., $\mu(X_j) \leq \mu(X_i)$). A tie can be resolved arbitrarily.
3. Party j subdivides X_j into two equal-population pieces, D_1 and D_2.
4. Party i chooses one of the pieces D_{k_j} for party j to redistrict.
5. Party j draws as many districts as they can within D_{k_j}, each of measure $\frac{1}{m}$.
6. Party i extends this to a complete m-partition of the state.

We show that each party has a strategy to guarantee its geometric target under this protocol. The full proof is technical and broken up into several lemmas; let us first give a high-level overview. In Step 1, each party k will submit as their X_k set a subset of $[0, 1]$ that they consider to be the "battleground" areas, where both parties have the same level of support, so either party could hope to win districts by gerrymandering. Formally, X_k is defined to be a set of maximal measure such that both parties have the same level of support over X_k according to party k's beliefs on the distribution of voters. Our first key observation (Lemmas 3 and 4) is that a best partition for party k is one that perfectly divides X_k into districts that k barely wins, while a worst partition is one that perfect divides X_k into districts that k barely loses, and whatever happens outside of X_k in these extreme cases is irrelevant. It follows that, in order for player k to get halfway from their worst possible utility to their best possible utility, thereby satisfying their geometric target, it suffices for player k to be granted control over redistricting (a specific) half of their X_k set. The cutter j thus satisfies their geometric target with an appropriate division in Step 3.

The difficult part of the proof lies in showing that the chooser i will be satisfied with at least one of the two choices in Step 4. To decide which piece is better, there are two different cases, depending on whether i believes they are a minority or a majority party. If i is a minority party, they use the partition of X_j to induce a partition of X_i into two equal pieces, and cede control over the piece in which they have less support, retaining control over the piece in which they have more support. It is then not too difficult to show that i will be able to meet its geometric target just from forming districts within the retained half of X_i.

The case where i is a majority party is more involved, since it may happen that, no matter which piece $D_{k_j} \subseteq X_j$ they cede to party j, it might be impossible for i to form enough districts from the remains to meet their geometric target. This is because both choices let party j form "packed" districts, in which party i wins by a large margin, wasting its advantage. However, when this happens, party i can respond by forming a packed district in $[0,1] \setminus D_{k_j}$ that party j wins for each packed district in D_{k_j} that party i wins. We argue that, for at least one of the two choices of k_j, party i will be left with majority over the remainder of the interval after forming these packed districts, so will be able to win all remaining districts. Since the wins in packed districts for each party exactly cancel each other out, this implies that party i meets their geometric target.

3.3. Running the GT Protocol on Real Geographic Data

Before proving Theorem 5, we take a brief digression to consider the practical validity of our approach, beyond the context of the state-cutting model. A sufficient condition for the existence of GT partitions in practice is that each step of the GT Protocol can be implemented in practice.

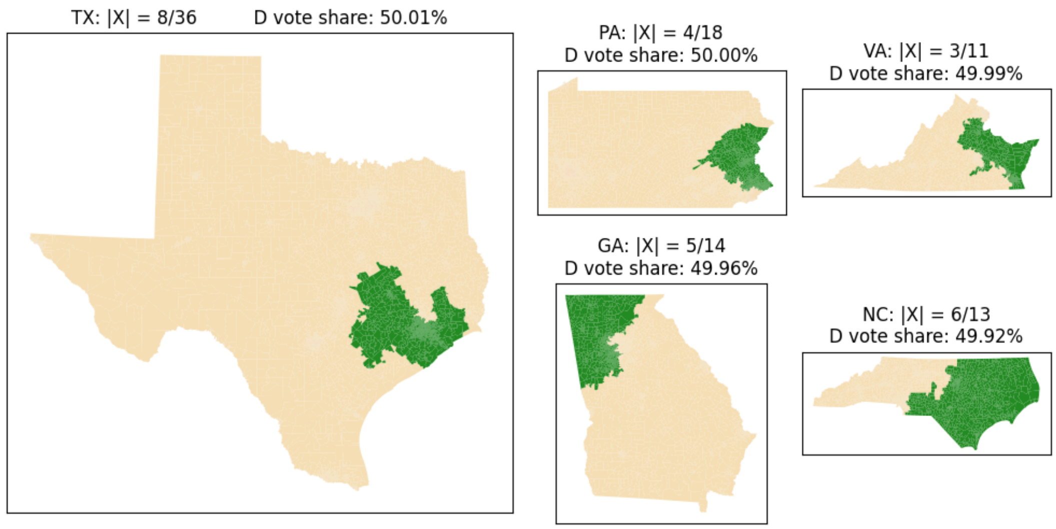

Figure 5 Connected, geographically compact X sets of the target size for all 6 of the states we study in this paper, except MA, where $|X| = 0$. Each region shown is the best of two random runs of our algorithm.

The GT Protocol relies crucially on Assumption 2. We have already seen a hypothetical example in the graph partitioning model where this fails: Theorem 3. If we try to execute the GT Protocol as stated on this instance, we run into trouble at Step 4, where party i must decide which part of X_j to cede control over. The proof of Theorem 5 gives a recipe for how party j should make this choice, looking at the intersections of subsets of X_i and X_j. In this instance, X_i and X_j violate Assumption 2: they are a series of horizontal strips and a series of vertical strips, and so their intersection is a jumbled set of disconnected vertices, which is what makes it impossible for party i to meet its geometric target. While such instances exist, we present empirical evidence suggesting that these too are the exception, not the norm.

The complete description of our algorithm and full set of empirical results are deferred to Appendix EC.6. Briefly, we first estimate how populous each X set should be using real election data from six US states (see more details on our data sources in Section 4). We then employ a heuristic optimization algorithm to attempt to find a compact set of the desired population, where both parties have nearly equal vote shares. We find that in all states, the X sets can be chosen to be connected and geographically compact; see Figure 5. This indicates that all of the steps of the GT Protocol could be faithfully executed to find a GT partition.

We note that, unlike in our state-cutting model, in practice a party would typically want a greater than 50% vote share to consider a district safe. We thus re-ran our algorithm several times for each state, targeting various other different vote shares, and we able to obtain similarly compact regions. These are presented in Appendix EC.6, along with histograms of Democratic vote shares of random regions of the target size.

3.4. Proof of Theorem 5

We begin by observing that it is possible to subdivide any district into two smaller districts of arbitrary sizes with the same fraction of party support as the original district. This is similar to the well-known "Austin Cut Procedure" from cake-cutting (Austin 1982).

LEMMA 1. *Given a voter distribution function v, a district D, and a real number $s \in [0,1]$, there exist districts D_1 and D_2 such that*

1. $D_1 \cup D_2 = D$,
2. $\mu(D_1 \cap D_2) = 0$,
3. $\mu(D_1) = s\mu(D)$, $\mu(D_2) = (1-s)\mu(D)$, *and*
4. $v(D_1) = sv(D)$, $v(D_2) = (1-s)v(D)$.

By iteratively applying Lemma 1, we obtain a more general form. The proof is completely straightforward, and hence omitted.

LEMMA 2. *Given a voter distribution function v, a district D, and $s \in \mathbb{R}_{>0} \cup \{\infty\}$, there exist districts $D_1, D_2, \ldots, D_{\lfloor 1/s \rfloor}$ such that*

1. *for all k, $D_k \subseteq D$,*
2. *for all $k_1 \neq k_2$, $\mu(D_{k_1} \cap D_{k_2}) = 0$,*
3. *for all k, $\mu(D_k) = s\mu(D)$, and*
4. *for all k, $v(D_k) = sv(D)$.*

For any $i, j \in N$, we say that j is a *minority party according to i* if $v_i^j([0,1]) \leq \frac{1}{2}$, and a *majority party according to i* if $v_i^j([0,1]) \geq \frac{1}{2}$. When $j = i$, we simply say i is a minority/majority party.

Note that this definition is merely with respect to the data of party i, so even if the inequalities are strict, it is still possible for both parties to be minority parties or both parties to be majority parties. Say that a district D is *competitive for* i if $v_i^j(D) = \frac{\mu(D)}{2}$ for some $j \in N$ (in which case it will clearly be true for all $j \in N$, since there are only two parties), and let

$$M_i := \{m\mu(D) \mid D \text{ is a competitive district for } i \text{ and } m\mu(D) \in \mathbb{Z}\}.$$

Since M_i is a nonempty set of integers that is bounded above (by m), it contains a maximum value. Let $m_i \in \mathbb{Z}_{\geq 0}$ be this maximum, and let X_i be one of the districts D attaining it, i.e., $m\mu(X_i) = m_i$. Note that m_i might be 0, in which case X_i is empty. Figure 6 shows the sets X_1 and X_2 for our running example (in this case they are both uniquely defined, up to adding sets of measure zero). Since $m = 10$, we have $m_1 = m\mu(X_1) = 7$ and $m_2 = m\mu(X_2) = 8$.

The foundation of our GT Protocol relies on the following key lemma, which will allow us to completely characterize geometric targets based on the sizes of X_1 and X_2:

Lemma 3. *For any $i, j \in N$, if j is a minority party according to i, then*

$$\min_{(P',T') \in \mathcal{P}(m)} u_i^j(P',T') = 0, \text{ and}$$

$$\max_{(P',T') \in \mathcal{P}(m)} u_i^j(P',T') = m_i.$$

The proof is deferred to Appendix EC.3. The main idea is that Lemma 2 implies the minority party may win zero districts by having their minority share exactly preserved in all of them. On the other hand, we prove using an application of the Intermediate Value Theorem (see Appendix EC.2) that their optimal partition sees them win m_i districts contained within X_i.

From this result we easily derive several consequences.

Lemma 4. *For any $i, j \in N$, if j is a majority party according to i, then*

$$\min_{(P',T') \in \mathcal{P}(m)} u_i^j(P',T') = m - m_i, \text{ and}$$

$$\max_{(P',T') \in \mathcal{P}(m)} u_i^j(P',T') = m.$$

Proof. Let j' denote the party that is not j. Note that j' must be a minority party according to i. For any m-partition (P',T') of $[0,1]$, $u_i^j(P',T')+u_i^{j'}(P',T')=m$. Therefore,

$$\begin{aligned}\min_{(P',T')\in\mathcal{P}(m)} u_i^j(P',T') &= \min_{(P',T')\in\mathcal{P}(m)} \left(m-u_i^{j'}(P',T')\right)\\ &= m-\max_{(P',T')\in\mathcal{P}(m)} u_i^{j'}(P',T')\\ &= m-m_i,\end{aligned}$$

where the final equality follows from Lemma 3 and the fact that j' is a minority party according to i. By the same reasoning, we analogously derive

$$\begin{aligned}\max_{(P',T')\in\mathcal{P}(m)} u_i^j(P',T') &= \max_{(P',T')\in\mathcal{P}(m)} \left(m-u_i^{j'}(P',T')\right)\\ &= m-\min_{(P',T')\in\mathcal{P}(m)} u_i^{j'}(P',T')\\ &= m.\end{aligned}$$

□

Lemma 5. *For any $i\in N$ and m-partition (P,T) of $[0,1]$, if party i is a minority party, then (P,T) satisfies the geometric target for i if and only if i wins at least $\lfloor \frac{m_i}{2}\rfloor$ districts under (P,T).*

Proof. This follows immediately from specializing $j:=i$ in Lemma 3, since the geometric target is for party i to win at least $\lfloor \frac{0+m_i}{2}\rfloor = \lfloor \frac{m_i}{2}\rfloor$ districts. □

Lemma 6. *For any $i\in N$ and m-partition (P,T) of $[0,1]$, if party i is a majority party, then (P,T) satisfies the geometric target for i if and only if i wins at least $m-\lceil \frac{m_i}{2}\rceil$ districts under (P,T).*

Proof. This follows from specializing $j:=i$ in Lemma 4, since the geometric target is for party i to win at least

$$\left\lfloor \frac{(m-m_i)+m}{2}\right\rfloor = \left\lfloor \frac{2m-m_i}{2}\right\rfloor = \left\lfloor m-\frac{m_i}{2}\right\rfloor = m+\left\lfloor \frac{-m_i}{2}\right\rfloor = m-\left\lceil \frac{m_i}{2}\right\rceil$$

districts. □

Lemma 7. *For any $i\in N$ and m-partition (P,T) of $[0,1]$, if party i wins at least $\lfloor \frac{m_i}{2}\rfloor$ competitive districts under (P,T), then (P,T) satisfies the geometric target for i.*

Proof. Let j denote the party that is not i. If i is a minority party, the result follows immediately from Lemma 5. If i is a majority party, then, by Lemma 6, the geometric target is for party i to win at least $m - \lceil \frac{m_i}{2} \rceil$ districts. Suppose toward a contradiction that (P,T) did not meet the geometric target for i, i.e., i wins strictly less than $m - \lceil \frac{m_i}{2} \rceil$ districts under (P,T). Let (P',T') be the m-partition of $[0,1]$ where $P' := P$ and $T'(D) := j$ for all $D \in P'$. With the new tie-breaking rule T', each of the $\lfloor \frac{m_i}{2} \rfloor$ competitive districts that party i won under (P,T) are instead won by party j according to i under (P',T'). Thus, party i wins $\lfloor \frac{m_i}{2} \rfloor$ fewer districts under (P',T'), which is strictly less than

$$\left(m - \left\lceil \frac{m_i}{2} \right\rceil\right) - \left\lfloor \frac{m_i}{2} \right\rfloor = m - m_i$$

districts in total. This contradicts the minimum value from Lemma 4. □

We are now ready to prove the main result. Below, we only consider the simpler case where the chooser party i is a minority party; the majority case is deferred to Appendix EC.4.

Proof of Theorem 5. We show that each party has a strategy to achieve its geometric target under the GT Protocol. For Step 1, we let X_1 and X_2 be defined as above, and then let $i,j \in \{1,2\}$ be defined as in Step 2, so that and $m_i \geq m_j$. For Step 3, we apply Lemma 1 to voter distribution function v_j^j, on district X_j, with $s := \frac{1}{2}$, obtaining districts D_1 and D_2 satisfying the four properties. See Figure 6 for an example of one valid choice of D_1 and D_2. Note that, for each $k \in \{1,2\}$, from property (3) of Lemma 1 we have

$$\mu(D_k) = \frac{\mu(X_j)}{2} = \frac{m_j}{2m}, \tag{1}$$

while from property (4), D_k is competitive for j since X_j is.

We claim that, for any choice of $k_j \in \{1,2\}$ in Step 4, it is possible for party j to create an m-partition of a subset of D_{k_j} in Step 5 such that, no matter how this partition is extended into an m-partition of $[0,1]$, the geometric target for party j is satisfied.

To prove this, we apply Lemma 2, to v_j^j, with $s := \frac{2}{m_j}$, to cut $\lfloor \frac{m_j}{2} \rfloor$ districts

$$P_{k_j} := \left\{ E_1, E_2, \ldots, E_{\lfloor \frac{m_j}{2} \rfloor} \right\}$$

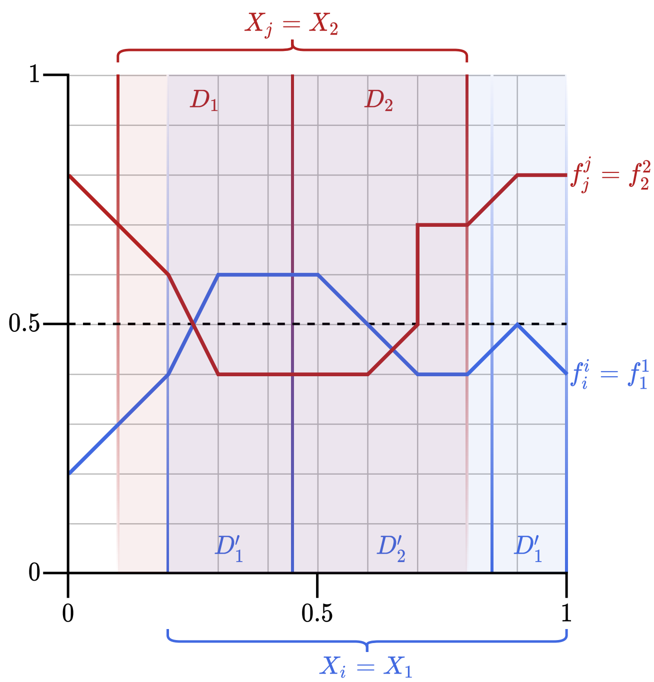

Figure 6 The same instance of the state-cutting problem from Figure 2, annotated with some of the sets described in the proof of Theorem 5. Note that $i = 1$ and $j = 2$ since $\mu(X_1) > \mu(X_2)$.

from D_{k_j}. From property (3) of Lemma 2, for each E_k district,

$$\begin{aligned}\mu(E_k) &= s\mu(D_{k_j}) \\ &= \frac{2}{m_j} \cdot \frac{m_j}{2m} \qquad \text{(from equation (1))} \\ &= \frac{1}{m},\end{aligned}$$

and, from property (4), each of these districts is competitive for j since D_{k_j} was. Thus, defining the tie-breaker over each E_k district by $T_{k_j}(E_k) := j$ ensures that party j wins all of these $\lfloor \frac{m_j}{2} \rfloor$ competitive districts under (P_{k_j}, T_{k_j}), so any extension of (P_{k_j}, T_{k_j}) satisfies the geometric target for j by Lemma 7.

It thus remains to establish that, for some $k_j \in \{1, 2\}$, we can extend (P_{k_j}, T_{k_j}) in Step 6 to an m-partition of $[0, 1]$ that satisfies the geometric target for party i. There are two cases, depending on whether party i is a minority or majority party (according to i).

Suppose i is a minority party (as in Figure 6). From equation (1) it follows that, for all $k \in \{1, 2\}$,

$$\mu(D_k \cap X_i) \le \mu(D_k) = \frac{m_j}{2m} \le \frac{m_i}{2m} = \frac{\mu(X_i)}{2}.$$

Therefore, it is possible to enlarge $D_1 \cap X_i$ and $D_2 \cap X_i$ into districts $D'_1, D'_2 \subseteq X_i$ that exactly partition X_i (ignoring overlapping endpoints of measure zero), both having equal measure

$$\mu(D'_k) = \frac{m_i}{2m} \tag{2}$$

(see Figure 6 for an example of a valid choice of D'_1 and D'_2). Since X_i is competitive for i,

$$0 = v_i^i(X_i) - \frac{\mu(X_i)}{2} = v_i^i(D'_1) + v_i^i(D'_2) - \frac{m_i}{2m} = \left(v_i^i(D'_1) - \frac{m_i}{4m}\right) + \left(v_i^i(D'_2) - \frac{m_i}{4m}\right).$$

Therefore, the two terms in parentheses cannot both be negative. Let $k_i \in \{1,2\}$ be such that

$$v_i^i(D'_{k_i}) \geq \frac{m_i}{4m}, \tag{3}$$

and let $k_j \in \{1,2\}$ be the other index, so $k_i \neq k_j$ (in Figure 6, $k_i = 1$ and $k_j = 2$). We construct an m-partition (P'_{k_i}, T'_{k_i}) by applying Lemma 2, to v_i^i, with $s := \frac{2}{m_i}$, to cut $\left\lfloor \frac{m_i}{2} \right\rfloor$ districts

$$P'_{k_i} := \left\{F_1, F_2, \ldots, F_{\left\lfloor \frac{m_i}{2} \right\rfloor}\right\}$$

from D'_{k_i}. According to property (3), each district F_k does indeed have the target size of

$$\begin{aligned}
\mu(F_k) &= s\mu(D'_{k_i}) \\
&= \frac{2}{m_i} \cdot \frac{m_i}{2m} \qquad \text{(from equation (2))} \\
&= \frac{1}{m}.
\end{aligned}$$

Furthermore, from property (4), each district F_k has party support

$$\begin{aligned}
v_i^i(F_k) &= s \cdot v_i^i(D'_{k_i}) \\
&= \frac{2}{m_i} \cdot v_i^i(D'_{k_i}) \\
&\geq \frac{2}{m_i} \cdot \frac{m_i}{4m} \qquad \text{(from inequality (3))} \\
&= \frac{1}{2m}.
\end{aligned}$$

We define the tie-breaker over each F_k district by $T'_{k_i}(F_k) := i$, ensuring that party i wins all of these $\left\lfloor \frac{m_i}{2} \right\rfloor$ districts. To form a GT partition for $[0,1]$, we take all districts and tie-breakers

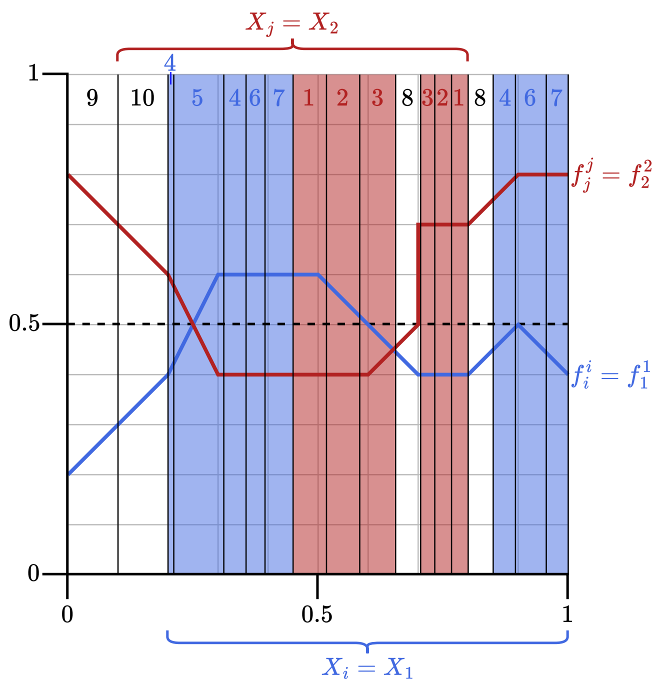

Figure 7 The final 10-partition meeting the geometric targets of both parties, with districts numbered in the order they are constructed in the proof of Theorem 5. The red districts 1-3 come from (P_{k_j}, T_{k_j}), so have ties broken in favor of party $j = 2$, while the blue districts 4-7 come from (P'_{k_i}, T'_{k_i}), so have ties broken in favor of party $i = 1$ (though in this case, it does not matter, since districts 4-7 are not competitive). The white districts 8-10 could be re-partitioned arbitrarily, and have ties broken in any way. Party $i = 1$ expects to win districts 1, 4, 5, 6, and 7, exceeding their geometric target of four districts, while party $j = 2$ expects to win all except district 5, exceeding their geometric target of six districts.

from (P'_{k_i}, T'_{k_i}) and (P_{k_j}, T_{k_j}) (which are necessarily disjoint since D'_{k_i} and D_{k_j} are), dividing the remainder of $[0,1]$ arbitrarily. Since party i is the minority party and wins at least $\lfloor \frac{m_i}{2} \rfloor$ districts, the geometric target for party i is satisfied by Lemma 5. Figure 7 shows the final 10-partition for our running example. The case where i is a majority party is handled in Appendix EC.4. $\square$

3.5. State-Cutting with Constraints

As an aside, we close our theoretical discussion by noting an extension involving additional constraints. As previously noted, our state-cutting model is designed to capture density constraints alone, not other geographic considerations like connectedness or compactness. While the topology of the unit interval is not the best vehicle for studying connectivity constraints in particular, it is natural to ask how simple constraints bounding the "complexity" of districts affect the feasibility of finding GT partitions. In Appendix EC.5 we show that GT partitions need *not* exist in the presence of this particular constraint.

4. GT Partitions in Practice

In this section we empirically investigate whether GT partitions exist in practice and what they look like. In the spirit of the *price of fairness* (Caragiannis et al. 2009, Bertsimas et al. 2011), we are particularly interested in the trade-off between satisfying the geometric target and various optimization objectives; that is, we investigate to what degree GT partitions are inferior to those that optimize traditional measures of quality.

A first challenge, though, is computation. Ideally, we would like to exactly optimize for the number of districts each party can win and use these optimal solutions to compute the geometric targets. Unfortunately, state-of-the-art machinery does not support exact optimization over the entire space of feasible partitions at the scale of real-world instances. We therefore rely on a heuristic evaluation of the extreme partitions; specifically, we use the GerryChain software developed by the Voting Rights Data Institute (Voting Rights Data Institute 2018) that efficiently implements the ReCom algorithm described in the previous section to generate thousands of valid partitions. Before a move to a new partition is accepted, it is verified that the new partition is contiguous and satisfies population equality to within 2% (with the exception of Virginia, where a bound of 5% is used). The precinct geometries and election data used in these experiments were prepared by the MGGG Redistricting Lab and are publicly available (Metric Geometry and Gerrymandering Group 2020). They include annotations for population, area, perimeter, and the number of Democratic and Republican votes cast in several recent elections.

It is worth noting that our positive theoretical results only hold in the state-cutting model with with exact population balance. However, we do not view our experimental setup, where we allow for 2-5% deviation, as a conceptual departure, as small population imbalances are allowed in practice and we have only relaxed the balance constraints enough to get computational tractability. It is difficult to find any valid redistricting plans at all with tight population constraints (Buchanan et al. 2025). We would therefore not suspect that this relaxation changes the space of potential maps in a way that effects our results. We return to this issue briefly in Section 5.

We generate 50 000 valid partitions (of which the first 1 000 are discarded) in six U.S. states: Georgia (GA), Massachusetts (MA), North Carolina (NC), Pennsylvania (PA), Texas (TX), and Virginia (VA). The relatively small number of steps in the Markov chain is due to the fact that we are using recombination moves (DeFord et al. 2021). If smaller, more local moves were used to traverse the space of partitions, several million would have been required (Metric Geometry and Gerrymandering Group 2018). At every partition we keep track of three metrics:

- The efficiency gap (Stephanopoulos and McGhee 2015), which measures the net difference in the fraction of each party's wasted votes. Every vote cast for the minority in a district is deemed to have been wasted, as are all votes for the majority above the threshold required to win the district.
- The number of competitive districts, defined to be those districts in which the majority party wins no more than 54% support.
- Compactness as measured by the Polsby-Popper (PP) score (Polsby and Popper 1991), computed as the ratio of the area of a district to the area of a circle with the same perimeter length.

Note that a smaller efficiency gap is better — a threshold of 8% is commonly accepted (Stephanopoulos and McGhee 2015) — while more compact districts have larger Polsby-Popper scores and we prefer a larger number of competitive districts.

Along with these metrics we compute the number of districts won by each party. This allows us to calculate the geometric targets and measure the price of fairness.

4.1. When Parties Agree About Voter Distributions

First, we consider the case where both parties agree about the distribution of voters. In this case we use the votes cast in the 2016 presidential election to evaluate the number of districts won by each party in every partition. Geometric targets are computed by taking the average (rounded down) of the minimum and maximum number of districts won by a party in any partition of the ensemble. In all of our experiments, we find that GT partitions exist. Table 1 reports the best observed values for each metric among GT partitions, as well as the optimal value observed among all partitions (when different).

	GA	**MA**	**NC**	**PA**	**TX**	**VA**
# Districts	14	9	13	18	36	11
Democratic vote share (%)	47.6	64.7	48.1	49.6	45.3	52.5
Democratic GT	4	9	5	7	15	7
Republican GT	9	0	8	11	21	4
Competitive districts	7	2	8	8	12 (13)	6
Efficiency gap (%)	6.0 (0*)	20.7	4.1 (0*)	5.4 (1.2)	0.1 (0*)	4.2 (0*)
Compactness (PP)	0.214	0.354	0.262	0.222 (0.225)	0.194 (0.2)	0.25

Table 1 For each state, its number of Congressional districts, the normalized Democratic vote share in the 2016 presidential election (calculated from official election results ignoring votes for third-party candidates), the Democratic and Republican geometric targets and, for each of three optimization objectives, the optimal value subject to satisfying the geometric target and the optimal value without this constraint (in parentheses, where different). Absolute efficiency gaps of 0^* do not exceed 0.05%.

We see in Table 1 that the cost of enforcing the geometric target as fairness constraint is very low. There is only one state (TX) in which this constraint leads to a decrease in the number of competitive districts compared to the maximum competitive districts observed. The decrease in compactness is never more than 3%. The increase in efficiency gap is larger (4-6%); however, we observe GT partitions meeting the recommended efficiency gap threshold of 8% in every state except MA, where meeting the threshold is impossible even without a fairness constraint.

We did not explicitly consider optimizing multiple objectives simultaneously; nevertheless, we observe several GT partitions that outperform the currently implemented partitions in these states on all three axes. Figure 8 shows two such GT partitions, one in Virginia and one in North Carolina.

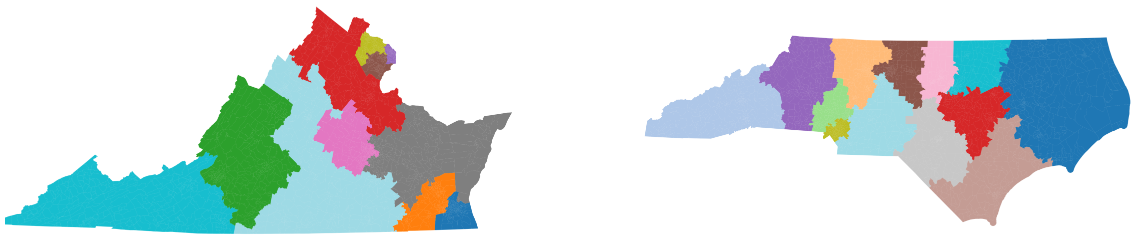

Figure 8 GT partitions in Virginia (left) and North Carolina (right) which outperform their implemented plans in terms of competitiveness, efficiency gap and compactness.

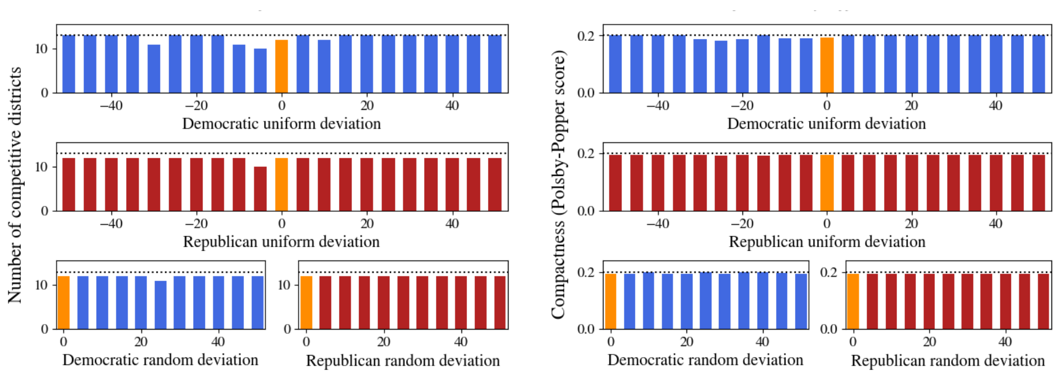

Figure 9 The largest number of competitive districts (left panel) and degree of compactness (right) of the best GT partitions observed in Texas when parties inflate or deflate their reported voter distribution by up to 50%. The black dotted line represents the maximum value observed among any partition. The color of the bar represents which party deviates. The golden bars report the experiment where neither party deviates, which is also recorded in the TX column of Table 1 (so in each panel, all four golden bars represent the same experiment).

The Virginian partition has three competitive districts (compared to two in their 2012 plan), an absolute efficiency gap of 6.6% (compared to 10.9%) and a compactness score of 0.185 (compared to 0.158). Similarly, the partition of North Carolina has three competitive districts (compared to 0 in their 2016 map), an efficiency gap of 7.1% (compared to 22.2%) and a Polsby-Popper score of 0.262 (compared to 0.252). The implemented plans are not only worse according to all three of the common optimization objectives we use as metrics, they also do not satisfy the geometric targets.

4.2. When Parties Disagree About Voter Distributions

A core strength of our theoretical result is that it does not require parties to agree on how voters will vote, as geometric targets can be guaranteed with respect to separate beliefs for each party. These divergent beliefs may be due to noisy data collection, polling errors or strategic manipulation.

To simulate such settings, we consistently let one of the parties report the true votes cast in the 2016 Presidential election, which we treat as the ground truth for the purpose of computing competitiveness and efficiency gaps. The other party's beliefs are allowed to deviate in several structured ways. First, we consider the case where the other party expects the votes to reflect the 2012 Presidential election (or, due to data availability in Georgia and Virginia, senate and congressional elections). Second, in an attempt to simulate possible strategic behavior, we consider what happens when the party uniformly under or over-reports their share of the votes in every region by $x\%$, for $x \in X = \{5, 10, \ldots, 50\}$. Finally, we consider the case where a party randomly inflates or deflates their share of the votes in each region (independently) by $y\%$, with $y \sim \text{Uniform}(-x, x), x \in X$.

As when parties agree about voter preferences, in all of our experiments, we find that GT partitions exist. Figure 9 compares the most competitive and compact GT partitions observed in Texas for each of the deviations we consider. In most of the scenarios, enforcing the geometric target led to the loss of at most one competitive district; the largest number of competitive districts lost was 3. In terms of compactness (measured by the Polsby-Popper score) the largest loss was when the Democratic party deflated their reported beliefs uniformly by 25%, leading to a GT partition with a compactness score of 0.183 compared to the unconstrained optimum of 0.200. The same trends held in the setting where the alternative voter distribution is from a different election. The effect of enforcing the geometric targets on competitiveness and compactness are similar in the other states, and we observed GT partitions meeting the efficiency gap threshold everywhere (with the obvious exception of MA). The full results from all experiments appear in EC.7.

Together these results tell a compelling story: not only is it easy to find GT partitions in practice, but restricting our search to GT partitions has little impact on the quality of the partition according to traditional optimization objectives.

5. Discussion

Our suggested redistricting approach relies on optimization subject to a fairness constraint. The fact that our fairness notion is readily satisfied in practice creates the opportunity to use it in isolation should optimization-based approaches prove impossible, either because of political objections or legislative difficulties. In such cases simply requiring that partitions meet the geometric target prevents the most extreme partisan outcomes yet allows legislators to retain much of the power and freedom that comes with the ability to decide where to draw district boundaries.

Through our state-cutting model, we have demonstrated how the powerful tools of fair division can be applied to the critically important problem of political redistricting. Previous theoretical investigations of fair redistricting have been stymied by modelling issues: geometric constraints are hard to justify and intractable to work with, so typically theorems are only proved in the trivial "geometry-free" model where there are no constraints whatsoever. We believe our state-cutting model strikes a useful balance between these extremes, distilling the key challenges of redistricting without explicitly considering geometry. It is a fertile ground on which fairness principles for redistricting can be rigorously tested. The intuitive geometric target criterion is one such principle, though we envision more to follow.

A shortcoming of our approach is the issue of computation. A specific problem is that using the minimum and maximum number of seats won by both parties across sampled partitions to compute the geometric targets does not necessarily lead to the true value: in theory, there could be more extreme partitions that were not observed. Regardless, we envision a process by which each party submits what it believes to be its best partition; the submitted partitions can then be used to compute the geometric target of each party. Under such a process, neither party would have a right to complain that it was disadvantaged in the computation of the geometric target.

The computation of GT partitions can also be incorporated into our theoretical model. We suspect that Robertson-Webb evaluation/cut queries (Robertson and Webb 1998) are insufficient to compute GT partitions, since it seems impossible to even compute the best and worst m-partitions for each party using this information, and thus it may be impossible to compute the X_i sets, which form the starting point of our protocol. Is there a richer query model under which it is possible to compute a GT partition using a finite number of queries?

There are a range of other directions in which our state-cutting model could be extended. The infeasibility in Theorem 1 shows that we cannot expect positive results when generalizing the state-cutting model to fully capture geography. Nevertheless, seeing how much geometry can be incorporated in the model may be an interesting direction for future work, including in two-dimensional variants of our state-cutting model. A separate question is whether we can still guarantee GT partitions exist when districts can deviate by some small percentage from their optimal size. This is surely true when parties agree on voter distributions (by the same proof of Theorem 4), but we leave the general question open. Another direction is to consider multiple parties. The first obstacle to extending beyond two parties is conceptual: it is unclear what the analogue of the geometric target is in that setting. We do not view this as a major limitation of our work though, as it is directly motivated by redistricting in the United States, which essentially has a two-party system.

These shortcomings notwithstanding, our results show that it is possible and practical to guarantee fairness even in a climate of extreme partisanship. We believe this insight could prove useful not just to academics, but also to state legislatures, courts, and independent redistricting commissions.

References

(1986) Davis v. Bandemer. 478 U.S. 109.

(2018) League of women voters of PA. v. com. of PA. 181 A.3d 1083.

A Akitaya H, Gonczi A, Souvaine DL, Tóth CD, Weighill T (2023) Reconfiguration of Polygonal Subdivisions via Recombination. Gørtz IL, Farach-Colton M, Puglisi SJ, Herman G, eds., *31st Annual European Symposium on Algorithms (ESA 2023)*, volume 274 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 6:1–6:16 (Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik), ISBN 978-3-95977-295-2, ISSN 1868-8969, URL `http://dx.doi.org/10.4230/LIPIcs.ESA.2023.6`.

Agarwal PK, Ko S, Munagala K, Taylor E (2022) Locally fair partitioning. *Thirty-Sixth AAAI Conference on Artificial Intelligence*, 4752–4759 (AAAI Press), URL `http://dx.doi.org/10.1609/AAAI.V36I5.20401`.

Allen D (2023) Opinion: Our democracy is menaced by two dragons. here's how to slay them. *The Washington Post* URL https://www.washingtonpost.com/opinions/2023/07/20/gerrymandering-electoral-college-solution-democracy/.

Arunachaleswaran ER, Barman S, Rathi N (2019) Fully polynomial-time approximation schemes for fair rent division. *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1994–2013.

Austin AK (1982) Sharing a cake. *The Mathematical Gazette* 66(437):212–215.

Aziz H, Mackenzie S (2016) A discrete and bounded envy-free cake cutting protocol for any number of agents. *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, 416–427.

Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Operations Research* 59(1):17–31.

Borodin A, Lev O, Shah N, Strangway T (2022) Little house (seat) on the prairie: Compactness, gerrymandering, and population distribution. *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*, 154–162, URL http://dx.doi.org/10.5555/3535850.3535869.

Bouie J (2023) A breathtaking contempt for the people of Wisconsin. *The New York Times* URL https://www.nytimes.com/2023/09/08/opinion/wisconsin-judge-impeachment-democracy.html.

Brams SJ (2020) Making partisan gerrymandering fair: One old and two new methods. *Social Science Quarterly* 101(1):68–72.

Brams SJ, Taylor AD (1996) *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press).

Buchanan A, Ezazipour S, Shahmizad M (2025) A widespread belief about county splits in political districting plans is wrong. *Election Law Journal: Rules, Politics, and Policy* .

Butler DE (1952) *The British General Election of 1951* (Macmillan).

Bycoffe A, Koeze E, Wasserman D, Wolfe J (2018) The atlas of redistricting. URL https://projects.fivethirtyeight.com/redistricting-maps/.

Cannon S, Duchin M, Randall D, Rule P (2022) Spanning tree methods for sampling graph partitions. arXiv preprint arXiv:2210.01401.

Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2009) The efficiency of fair division. *Proceedings of the 5th Workshop on Internet and Network Economics (WINE)*, 475–482.

Charikar M, Liu P, Liu T, Vuong TD (2022) On the complexity of sampling redistricting plans. arXiv preprint arXiv:2206.04883.

De Silva J, Gales B, Kagy B, Offner D (2018) An analysis of a fair division protocol for drawing legislative districts. arXiv:1811.05705.

DeFord D, Duchin M, Solomon J (2021) Recombination: A Family of Markov Chains for Redistricting. *Harvard Data Science Review* 3(1), URL `https://hdsr.mitpress.mit.edu/pub/1ds8ptxu`.

Dehghani S, Farhadi A, Hajiaghayi M, Yami H (2018) Envy-free chore division for an arbitrary number of agents. *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2564–2583.

Donnay C, Kahle M (2025) Asymptotics of redistricting the $n \times n$ grid. *The American Mathematical Monthly* 132(9):856–866.

Duchin M (2018) Geometry v. gerrymandering. *Scientific American* 319(5):48–53.

Duchin M, Gladkova T, Henninger-Voss E, Klingensmith B, Newman H, Wheelen H (2019) Locating the representational baseline: Republicans in massachusetts. *Election Law Journal: Rules, Politics, and Policy* 18(4):388–401.

Edmonds J, Pruhs K (2006) Cake cutting really is not a piece of cake. *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 271–278.

Frieze A, Pegden W (2023) Subexponential mixing for partition chains on grid-like graphs. *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 3317–3329, URL `http://dx.doi.org/10.1137/1.9781611977554.ch127`.

Gal Y, Mash M, Procaccia AD, Zick Y (2017) Which is the fairest (rent division) of them all? *Journal of the ACM* 64(6):article 39.

Garfinkel RS, Nemhauser GL (1970) Optimal political districting by implicit enumeration techniques. *Management Science* 16(8):495–508.

Gelman A, King G (1994) A unified method of evaluating electoral systems and redistricting plans. *American Journal of Political Science* 38(2):514–554.

Gomberg A, Pancs R, Sharma T (2023a) Electoral maldistricting. *International Economic Review* 64(3):1223–1263.

Gomberg A, Pancs R, Sharma T (2023b) Padding and pruning: Gerrymandering under turnout heterogeneity. Available at SSRN: `https://ssrn.com/abstract=4436682`.

Grofman B, King G (2007) The future of partisan symmetry as a judicial test for partisan gerrymandering after lulac v. perry. *Election Law Journal* 6(1):2–35.

Gurnee W, Shmoys DB (2021) Fairmandering: A column generation heuristic for fairness-optimized political districting. *SIAM Conference on Applied and Computational Discrete Algorithms (ACDA21)*, 88–99 (SIAM).

Hess SW, Weaver JB, Siegfeldt HJ, Whelan JN, Zitlau PA (1965) Nonpartisan political redistricting by computer. *Operations Research* 13:998–1006.

Jackman S (1994) Measuring electoral bias: Australia, 1949–93. *British Journal of Political Science* 24(3):319–357.

King G (1989) Representation through legislative redistricting: A stochastic model. *American Journal of Political Science* 33(4):787–824.

Landau Z, Reid O, Yershov I (2009) A fair division solution to the problem of redistricting. *Social Choice and Welfare* 32(3):479–492.

Landau Z, Su FE (2014) Fair division and redistricting. *The Mathematics of Decisions, Elections, and Games*, 17–36 (American Mathematical Society).

Mehrotra A, Johnson EL, Nemhauser GL (1998) An optimization based heuristic for political districting. *Management Science* 44(8):1100–1114.

Metric Geometry and Gerrymandering Group (2018) Comparison of districting plans for the Virginia House of Delegates. Technical report.

Metric Geometry and Gerrymandering Group (2020) mggg-states: Open collection of precincts shapefiles for U.S. states. URL `https://github.com/mggg-states`.

Moulin H (2003) *Fair Division and Collective Welfare* (MIT Press).

Nagle JF (2017) How competitive should a fair single member districting plan be? *Election Law Journal* 16(1):196–207.

Niemi RD, Deegan J (1978) A theory of political districting. *American Political Science Review* 72(4):1304–1323.

Oehrlein J, Haunert J (2017) A cutting-plane method for contiguity-constrained spatial aggregation. *Journal of Spatial Information Science* 2017(15):89–120.

Pancs R, Sharma T (2025) One man one vote.

Pegden W, Procaccia AD, Yu D (2017) A partisan districting protocol with provably nonpartisan outcomes. arXiv preprint arXiv:1710.08781.

Polsby D, Popper R (1991) The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale Law & Policy Rev.* 9:301.

Procaccia AD (2013) Cake cutting: Not just child's play. *Communications of the ACM* 56(7):78–87.

Procaccia AD (2019) Axioms should explain solutions. Laslier JF, Moulin H, Sanver R, Zwicker WS, eds., *The Future of Economic Design* (Springer).

Procaccia AD, Tucker-Foltz J (2022) Compact redistricting plans have many spanning trees. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* .

Robertson JM, Webb WA (1998) *Cake Cutting Algorithms: Be Fair If You Can* (A. K. Peters).

Schutzman Z (2020) Trade-offs in fair redistricting. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 159–165.

Shirabe T (2005) A model of contiguity for spatial unit allocation. *Geographical Analysis* 37(1):2–16.

Shirabe T (2009) Districting modeling with exact contiguity constraints. *Environment and Planning B: Planning and Design* 36(6):1053–1066.

Stephanopoulos NO, McGhee EM (2015) Partisan gerrymandering and the efficiency gap. *University of Chicago Law Review* 82:831.

Swamy R, King DM, Jacobson SH (2023) Multiobjective optimization for politically fair districting: A scalable multilevel approach. *Operations Research* 71(2):536–562.

Tapp K (2021) Spanning tree bounds for grid graphs. arXiv preprint arXiv:2109.05987.

Tucker-Foltz J (2018) A cut-and-choose mechanism to prevent gerrymandering. arXiv preprint arXiv:1802.08351.

Tucker-Foltz J (2023) Locked polyomino tilings. arXiv preprint arXiv:2307.15996.

Validi H, Buchanan A, Lykhovyd E (2022) Imposing contiguity constraints in political districting models. *Operations Research* 70(2):867–892.

Voting Rights Data Institute (2018) Gerrychain. URL `https://github.com/mggg/GerryChain`.

Proofs omitted from main body

EC.1. Proof of Lemma 1

Let f denote the density function of v. Without loss of generality, we may assume that $D=[0,t]$ where $t=\mu(D)$, for otherwise we could simply rearrange the finite number of intervals comprising D so that this is the case and adapt the proof accordingly. Define functions $g:[0,2t]\to[0,1]$ by

$$g(x):=\begin{cases} f(x) & \text{if } x\le t \\ f(x-t) & \text{if } x>t \end{cases}$$

and $h:[0,t]\to[0,1]$ by

$$h(x):=\int_x^{x+st} g(y)dy.$$

Intuitively, for any $x\in[0,t]$, $h(x)$ is the value of a piece of measure st that begins at x, wrapping around if necessary. Observe that the average value of h over $[0,t]$ is

$$\begin{aligned}
\frac{1}{t}\int_0^t h(x)dx &= \frac{1}{t}\int_0^t\int_x^{x+st} g(y)dydx \\
&= \frac{1}{t}\int_0^t\int_0^{st} g(x+y)dydx \\
&= \frac{1}{t}\int_0^{st}\left(\int_0^t g(x+y)dx\right)dy \\
&= \frac{1}{t}\int_0^{st}\left(\int_y^{t+y} g(x)dx\right)dy \\
&= \frac{1}{t}\int_0^{st}\left(\int_y^t g(x)dx+\int_t^{t+y} g(x)dx\right)dy \\
&= \frac{1}{t}\int_0^{st}\left(\int_y^t g(x)dx+\int_0^{y} g(x+t)dx\right)dy \\
&= \frac{1}{t}\int_0^{st}\left(\int_y^t f(x)dx+\int_0^{y} f(x)dx\right)dy \\
&= \frac{1}{t}\int_0^{st}\left(\int_0^t f(x)dx\right)dy \\
&= \frac{1}{t}(st)\int_0^t f(x)dx \\
&= sv(D).
\end{aligned}$$

Since h is clearly continuous, by the intermediate value theorem there must exist some $x^* \in [0,t]$ at which h attains its average value. If $x^* + st \le t$, then we define

$$D_1 := [x^*, x^* + st].$$

In this case,

$$\mu(D_1) = st = s\mu(D)$$

and

$$v(D_1) = \int_{x^*}^{x^*+st} f(y)dy = \int_{x^*}^{x^*+st} g(y)dy = h(x^*) = sv(D).$$

If $x^* + st > t$, we instead define

$$D_1 := [x^*, t] \cup [0, x^* + st - t].$$

Note that these intervals are both contained within $D = [0,t]$ and are non-overlapping since $s \le 1$ (except possibly at the point x^* in the case where $s = 1$). Therefore,

$$\mu(D_1) = (t - x^*) + (x^* + st - t) = st = s\mu(D)$$

and

$$\begin{aligned} v(D_1) &= \int_{x^*}^{t} f(y)dy + \int_{0}^{x^*+st-t} f(y)dy \\ &= \int_{x^*}^{t} f(y)dy + \int_{t}^{x^*+st} f(y-t)dy \\ &= \int_{x^*}^{t} g(y)dy + \int_{t}^{x^*+st} g(y)dy \\ &= \int_{x^*}^{x^*+st} g(y)dy \\ &= h(x^*) \\ &= sv(D). \end{aligned}$$

Thus, in either case, we have found a district D_1 satisfying properties (3) and (4). Letting

$$D_2 := \overline{D \setminus D_1}$$

(the closure of $D \setminus D_1$), properties (1) and (2) are automatically satisfied. Furthermore,

$$\mu(D_2) = \mu(D) - \mu(D_1) = \mu(D) - s\mu(D) = (1-s)\mu(D)$$

and

$$v(D_2) = v(D) - v(D_1) = v(D) - sv(D) = (1-s)v(D),$$

so D_2 satisfies properties (3) and (4) as well. □

EC.2. An application of the Intermediate Value Theorem

We will use the following lemma multiple times in the proofs that follow.

LEMMA EC.1. *For any $i, j \in N$, let Y be a district such that one of*

$$v_i^j(Y) - \frac{\mu(Y)}{2} \quad \text{and} \quad v_i^j([0,1]) - \frac{1}{2}$$

is ≥ 0 and the other is ≤ 0. Then $\mu(Y) < \frac{m_i+1}{m}$.

Proof. Suppose towards contradiction $\mu(Y) \geq \frac{m_i+1}{m}$. Define a function $g : [0,1] \to [-\frac{1}{2}, \frac{1}{2}]$ by

$$g(t) := v_i^j(Y \cup [0,t]) - \frac{\mu(Y \cup [0,t])}{2}.$$

Clearly, g is continuous. Furthermore,

$$g(0) = v_i^j(Y) - \frac{\mu(Y)}{2},$$
$$g(1) = v_i^j([0,1]) - \frac{\mu([0,1])}{2} = v_i^j([0,1]) - \frac{1}{2}.$$

By assumption, one of these terms must be ≥ 0 and the other ≤ 0. By the intermediate value theorem, there exists $t^* \in [0,1]$ such that $g(t^*) = 0$. Letting $D := Y \cup [0, t^*]$, we must have that

$$v_i^j(D) = g(t^*) + \frac{\mu(D)}{2} = \frac{\mu(D)}{2},$$

i.e., D is competitive for i. Since $Y \subseteq D$, $\mu(D) \geq \mu(Y) \geq \frac{m_i+1}{m}$. Thus, we may apply Lemma 1 to voter distribution function v_i^j, with $s := \frac{m_i+1}{m\mu(D)} \in [0,1]$, to cut out a district $D_1 \subseteq D$ of measure

$$\mu(D_1) = \frac{m_i+1}{m\mu(D)} \cdot \mu(D) = \frac{m_i+1}{m}.$$

Furthermore, observe that, since D is competitive for i, it follows from property (4) of Lemma 1 that D_1 is competitive for i:

$$\begin{aligned}
v_i^j(D_1) &= s \cdot v_i^j(D) \\
&= s \cdot \frac{\mu(D)}{2} \qquad \text{(because } D \text{ is competitive for } i\text{)} \\
&= \frac{m_i + 1}{m\mu(D)} \cdot \frac{\mu(D)}{2} \\
&= \frac{m_i + 1}{2m} \\
&= \frac{\mu(D_1)}{2}.
\end{aligned}$$

This proves $m_i + 1 \in M_i$, contradicting the definition of m_i as the maximum element of M_i. $\square$

EC.3. Proof of Lemma 3

Let j' denote the party that is not j. To prove the first equation, we apply Lemma 2 to v_i^j, with $s := \frac{1}{m}$, to divide $[0,1]$ into m districts $D_1, D_2, \ldots, D_m$ of equal size $\frac{1}{m}$. In each district D_k, from property (4) of Lemma 2 and the fact that j is a minority party according to i,

$$v_i^j(D_k) = \frac{1}{m} v_i^j([0,1]) \leq \frac{1}{2m}.$$

Therefore, if we break ties in favor of party j', party j will win none of these districts. Formally, letting $P' := \{D_1, D_2, \ldots, D_m\}$ and $T'(D_k) := j'$ for each $k \in [m]$, we have that $u_i^j(P', T') = 0$, proving the first equation.

To prove the second equation, we apply Lemma 2 to v_i^j, with $s := \frac{1}{m_i}$, to divide X_i into m_i districts $D_1, D_2, \ldots, D_{m_i}$. Note that, by property (3) of Lemma 1, each district D_k has size

$$\mu(D_k) = s \cdot \mu(X_i) = \frac{1}{m_i} \cdot \frac{m_i}{m} = \frac{1}{m}.$$

Furthermore, since X_i is competitive for i, it follows from property (4) of Lemma 2 that each D_k is competitive for i. Let P' consist of $D_1, D_2, \ldots, D_{m_i}$, along with an arbitrary division of $\overline{[0,1] \setminus X_i}$ (the closure of the complement of X_i) into $m - m_i$ districts, and let $T'(D_k) := j$ for each $k \in [m_i]$, with an arbitrary tie-breaking choice for all of the other districts. Since the D_k districts are competitive and ties are broken in favor of party j, it follows that party j will win each of them according to i. Therefore, $u_i^j(P',T') \geq m_i$, which proves that $\max_{(P',T') \in \mathcal{P}(m)} u_i^j(P',T') \geq m_i$.

To prove the other direction, suppose toward a contradiction that, for some m-partition (P',T') of $[0,1]$, $u_i^j(P',T') \geq m_i + 1$. Let $Y \subseteq [0,1]$ be the union of all districts won by j according to i under (P',T'). Since there are at least $m_i + 1$ such districts, each of measure $\frac{1}{m}$, we have

$$\mu(Y) \geq \frac{m_i + 1}{m}. \tag{EC.1}$$

However, $v_i^j(Y) - \frac{\mu(Y)}{2} \geq 0$ since party j wins each of the districts comprising Y according to i, and $v_i^j([0,1]) - \frac{1}{2} \leq 0$ since party j is a minority party according to i. By Lemma EC.1, we have $\mu(Y) < \frac{m_i+1}{m}$, contradicting inequality (EC.1). $\square$

EC.4. Proof of Theorem 5, majority case

In the case where i is a majority party, we first extend (P_1,T_1) and (P_2,T_2) by adding disjoint districts of size $\frac{1}{m}$ to (P_1,T_1) and (P_2,T_2), in alternation, until the total measure covered by $P_1 \cup P_2$ is exactly $\mu(X_i) = \frac{m_i}{m}$ (this is possible since $m_j \leq m_i$). Call the resulting partitions (P_1',T_1') and (P_2',T_2'). For each $k \in \{1,2\}$, let $A_k, B_k \subseteq [0,1]$ be comprised of all districts that party i wins/loses under (P_k',T_k'), respectively. Note that A_1, A_2, B_1, and B_2 are pairwise disjoint, have measures that are integer multiples of $\frac{1}{m}$, and for each $k \in \{1,2\}$,

$$\mu(A_k \cup B_k) \leq \frac{\left\lceil \frac{m_i}{2} \right\rceil}{m}. \tag{EC.2}$$

(This follows since both partitions started with the same number of districts and alternately grew one district at a time until reaching m_i districts, so the maximum number of districts either partition could have at the end is $\left\lceil \frac{m_i}{2} \right\rceil$.) Let C be the remaining part of the interval,

$$C := \overline{[0,1] \setminus (A_1 \cup A_2 \cup B_1 \cup B_2)}.$$

There are a few different sub-cases to consider, depending on the advantage of party i in each of these five districts. First suppose that, for some $k_j \in \{1,2\}$,

$$v_i^i(A_{k_j} \cup B_{k_j}) \leq \frac{\mu(A_{k_j} \cup B_{k_j})}{2}. \tag{EC.3}$$

Then it must be the case that

$$v_i^i(\overline{[0,1]\setminus(A_{k_j}\cup B_{k_j})}) \geq \frac{\mu(\overline{[0,1]\setminus(A_{k_j}\cup B_{k_j})})}{2}, \tag{EC.4}$$

for otherwise, summing (EC.3) with the negation of (EC.4), we would have that $v_i^i([0,1]) < \frac{1}{2}$, contradicting the assumption that i is a majority party. We apply Lemma 2 to divide $\overline{[0,1]\setminus(A_{k_j}\cup B_{k_j})}$ into $m - \lceil \frac{m_i}{2} \rceil$ districts of size $\frac{1}{m}$. Property (4) of Lemma 2 and inequality (EC.4) imply that party i will win all of these districts (as long as we break ties in favor of i). Thus, using these disjoint districts to extend (P'_{k_j}, T'_{k_j}) (which is itself an extension of (P_{k_j}, T_{k_j})), by Lemma 6, we have met the geometric target for party i.

Now suppose instead that, for all $k \in \{1,2\}$,

$$v_i^i(A_k \cup B_k) \geq \frac{\mu(A_k \cup B_k)}{2}. \tag{EC.5}$$

If, in addition, we have

$$v_i^i(C) \geq \frac{\mu(C)}{2},$$

then inequality (EC.4) clearly still holds for either choice of k_j, so the same argument goes through. Thus, assume

$$v_i^i(C) \leq \frac{\mu(C)}{2}. \tag{EC.6}$$

We claim that

$$\mu(C) \leq \mu(A_1) + \mu(A_2). \tag{EC.7}$$

Suppose toward a contradiction that (EC.7) did not hold. Since all three measures are integer multiples of $\frac{1}{m}$, this means that

$$\mu(C) \geq \mu(A_1) + \mu(A_2) + \frac{1}{m}. \tag{EC.8}$$

We proceed similarly as in the last part of the proof of Lemma 3. Letting $Y := B_1 \cup B_2 \cup C$, we have

$$\mu(Y) = \mu(B_1) + \mu(B_2) + \mu(C)$$

$$\begin{aligned}
&\geq \mu(A_1)+\mu(A_2)+\mu(B_1)+\mu(B_2)+\frac{1}{m} \qquad \text{(from inequality (EC.8))}\\
&= \frac{m_i}{m}+\frac{1}{m} \qquad \text{(by the definitions of } (P_1',T_1') \text{ and } (P_2',T_2'))\\
&= \frac{m_i+1}{m}. \qquad\qquad \text{(EC.9)}
\end{aligned}$$

However,

$$v_i^i(Y)-\frac{\mu(Y)}{2}\leq 0$$

from inequality (EC.6) and the fact that party i loses all districts in B_1 and B_2, and

$$v_i^i([0,1])-\frac{1}{2}\geq 0$$

since party i is a majority party. Therefore, by Lemma EC.1, we have

$$\mu(Y)<\frac{m_i+1}{m},$$

contradicting inequality (EC.9).

Thus, we have shown that inequality (EC.7) holds. It is therefore possible to subdivide C into two districts C_1 and C_2 such that, for each $k\in\{1,2\}$,

$$\mu(C_k)\leq \mu(A_k). \qquad \text{(EC.10)}$$

Since i is a majority party, and A_1, B_1, C_1, A_2, B_2, and C_2 form a partition of $[0,1]$ into districts that only overlap at endpoints,

$$\begin{aligned}
0&\leq v_i^i([0,1])-\frac{\mu([0,1])}{2}\\
&=\left(v_i^i(A_1\cup B_1\cup C_1)-\frac{\mu(A_1\cup B_1\cup C_1)}{2}\right)+\left(v_i^i(A_2\cup B_2\cup C_2)-\frac{\mu(A_2\cup B_2\cup C_2)}{2}\right),
\end{aligned}$$

so the two terms in parentheses cannot both be negative. Let $k_i\in\{1,2\}$ be such that

$$v_i^i(A_{k_i}\cup B_{k_i}\cup C_{k_i})\geq \frac{\mu(A_{k_i}\cup B_{k_i}\cup C_{k_i})}{2}$$

and let $k_j\in\{1,2\}$ be the other index, so $k_i\neq k_j$. As was done in the case where party i was the minority party, we extend (P_{k_j}',T_{k_j}') (which is itself an extension of (P_{k_j},T_{k_j})) by applying Lemma 2 to v_i^i with

$$s:=\frac{1}{m\mu(A_{k_i}\cup B_{k_i}\cup C_{k_i})}$$

to cut $\lfloor m\mu(A_{k_i} \cup B_{k_i} \cup C_{k_i})\rfloor$ districts from $A_{k_i} \cup B_{k_i} \cup C_{k_i}$, which we can ensure are all won by party i by breaking ties in favor of party i. Note that these districts clearly have the target size $\frac{1}{m}$ from property (3) of Lemma 2. The remainder of $[0,1]$ can be partitioned arbitrarily; denote by (P,T) the resulting m-partition of $[0,1]$. Recall that party i also wins all $m\mu(A_{k_j})$ districts from A_{k_j}. Thus, the total number of districts they win is

$$\begin{aligned}
u_i^i(P,T) &\geq \lfloor m\mu(A_{k_i} \cup B_{k_i} \cup C_{k_i})\rfloor + m\mu(A_{k_j}) \\
&= m\mu(A_{k_i} \cup B_{k_i} \cup C_{k_i}) + m\mu(A_{k_j}) \qquad \text{(since } \mu(A_{k_i} \cup B_{k_i} \cup C_{k_i}) \text{ is a multiple of } 1/m\text{)} \\
&= m\left(\mu(A_{k_i}) + \mu(B_{k_i}) + \mu(C_{k_i}) + \mu(A_{k_j})\right) \\
&\geq m\left(\mu(A_{k_i}) + \mu(B_{k_i}) + \mu(C_{k_i}) + \mu(C_{k_j})\right) \qquad \text{(from inequality (EC.10))} \\
&= m\left(1 - \mu(A_{k_j}) - \mu(B_{k_j})\right) \\
&= m - m\mu(A_{k_j} \cup B_{k_j}) \\
&\geq m - \left\lceil \frac{m_i}{2} \right\rceil \qquad \text{(from inequality (EC.2)).}
\end{aligned}$$

Hence, by Lemma 6, the geometric target of party i is satisfied. □

EC.5. State-Cutting with a Bounded Number of Intervals Per District

The state-cutting problem becomes uninteresting if we require that each part P_i of a partition is connected in the standard topological sense. There is only one such partition of $[0,1]$ into equal-sized pieces, which necessarily is both the best and worst partition for each party. This partition is therefore trivially a GT partition.

Now suppose we require each district to be a union of a bounded number of intervals. Formally, for positive integers m and t, we let $\mathcal{P}(m,t)$ denote the set of m-partitions into districts that are each a union of at most t intervals. Building off of our previous definition of a GT partition, we say that a *t-bounded GT partition* is an m-partition $(P,T) \in \mathcal{P}(m,t)$ such that, for all $i \in N$, the *t-bounded geometric target for party i* is satisfied:

$$u_i^i(P,T) \geq \left\lfloor \frac{1}{2}\left(\min_{\substack{(P',T')\\ \in\mathcal{P}(m,t)}} u_i^i(P',T') + \max_{\substack{(P',T')\\ \in\mathcal{P}(m,t)}} u_i^i(P',T')\right)\right\rfloor.$$

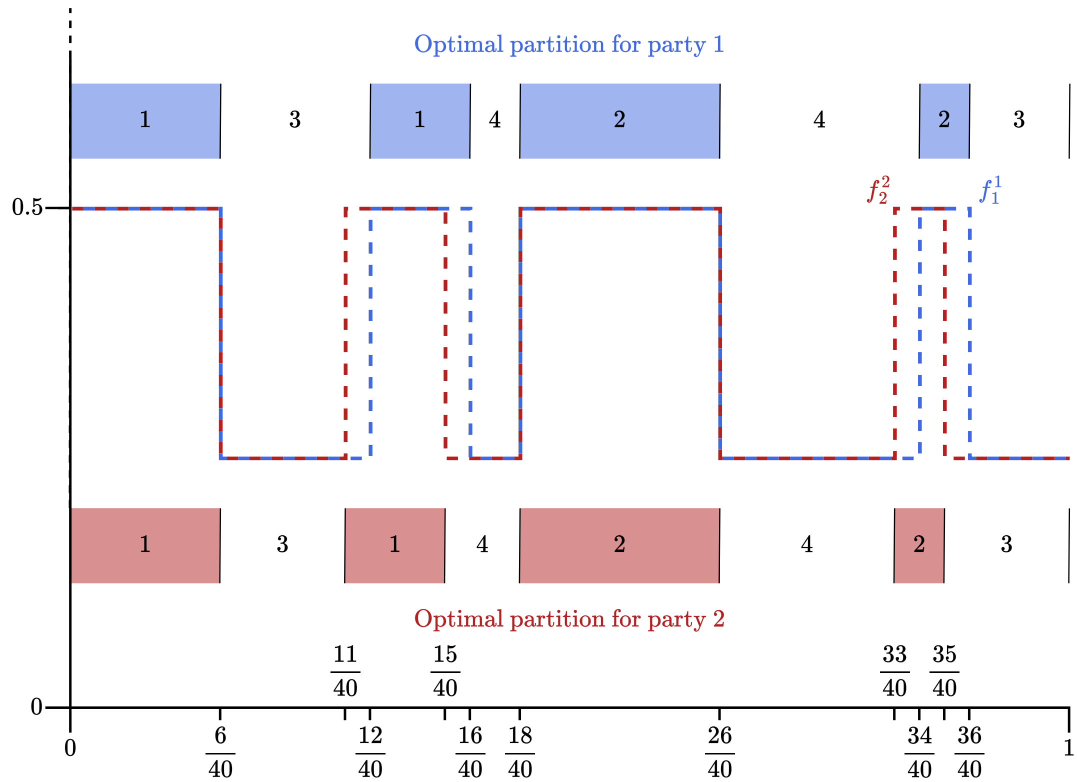

Figure EC.1 An instance of the state-cutting problem, with $m = 4$, in which and there is no 2-bounded GT partition. The density functions f_1^1 and f_2^2 are shown in blue and red, respectively, as dashed lines. Note that these functions coincide except in 4 small intervals. Above and below are shown the optimal partitions for each party, which award 2 districts (highlighted respectively in blue and red).

THEOREM EC.1. *Given any instance of the state-cutting problem in which $f_1^1 \equiv f_2^1$, a t-bounded GT partition exists for any t. When $f_1^1 \not\equiv f_2^1$, t-bounded GT partitions may not exist even for $t = 2$.*

Proof. The positive claim follows from the following observation about the proof of Theorem 4: The GT partition (P, T) constructed in the proof is obtained by rearranging and merging intervals from one of the two extreme partitions (P_1, T_1) or (P_2, T_2). Thus, each district in (P, T) has weakly fewer intervals than in some (P_i, T_i), so if $(P_i, T_i) \in \mathcal{P}(m, t)$ for a given value of t, it follows that $(P, T) \in \mathcal{P}(m, t)$ as well. Thus, (P, T) is a t-bounded GT partition.

For the negative claim, consider the counterexample depicted in Figure EC.1. Each party is a minority party according to itself, and clearly both parties expect to win 0 districts in their respective worst 4-partitions; for example, one of these worst partitions is

$$\left\{\left[0,\frac{1}{4}\right],\left[\frac{1}{4},\frac{1}{2}\right],\left[\frac{1}{2},\frac{3}{4}\right],\left[\frac{3}{4},1\right]\right\}.$$

The unique optimal partition for party 1 (blue) is

$$\left\{\left[0,\frac{6}{40}\right]\cup\left[\frac{12}{40},\frac{16}{40}\right],\left[\frac{18}{40},\frac{26}{40}\right]\cup\left[\frac{34}{40},\frac{36}{40}\right],\left[\frac{6}{40},\frac{12}{40}\right]\cup\left[\frac{36}{40},1\right],\left[\frac{16}{40},\frac{18}{40}\right]\cup\left[\frac{26}{40},\frac{34}{40}\right]\right\}.$$

If we break ties in favor of party 1, they will win the first 2 districts. Similarly, the unique optimal partition for party 2 (red) is

$$\left\{\left[0,\frac{6}{40}\right]\cup\left[\frac{11}{40},\frac{15}{40}\right],\left[\frac{18}{40},\frac{26}{40}\right]\cup\left[\frac{33}{40},\frac{35}{40}\right],\left[\frac{6}{40},\frac{11}{40}\right]\cup\left[\frac{35}{40},1\right],\left[\frac{15}{40},\frac{18}{40}\right]\cup\left[\frac{26}{40},\frac{33}{40}\right]\right\}.$$

If we break ties in favor of party 2, they will win the first 2 districts. Thus, to simultaneously meet the geometric targets of both parties, they must each win 1 of the 4 districts according to their respective valuation functions. We claim that no $(P,T)\in\mathcal{P}(4,2)$ has this property.

To meet the geometric target for party 1, we must form a district $D^1\in P$ of measure $\frac{1}{m}=\frac{10}{40}$ comprised of 2 closed intervals, contained within the subset of $[0,1]$ where v_1^1 has equal density for each party, breaking ties in favor of party 1. This subset is precisely

$$X_1=\left[0,\frac{6}{40}\right]\cup\left[\frac{11}{40},\frac{15}{40}\right]\cup\left[\frac{18}{40},\frac{26}{40}\right]\cup\left[\frac{33}{40},\frac{35}{40}\right].$$

Let us call the four intervals above respectively D_6^1, D_4^1, D_8^1, and D_2^1, so that D_k^1 has measure $\frac{k}{40}$. Similarly, to meet the geometric target for party 2, we must form a district $D^2\in P$ of measure $\frac{1}{m}=\frac{10}{40}$ comprised of 2 closed intervals, contained within

$$X_2=\left[0,\frac{6}{40}\right]\cup\left[\frac{12}{40},\frac{16}{40}\right]\cup\left[\frac{18}{40},\frac{26}{40}\right]\cup\left[\frac{34}{40},\frac{36}{40}\right].$$

We analogously call these four intervals D_6^2, D_4^2, D_8^2, and D_2^2. Since we have to break ties differently for each district, we must have $D^1\neq D^2$. (Note: if we raise each valuation function $v\in\{v_1^1,v_2^2\}$ by a sufficiently small amount ε over all places where v is strictly greater than the other valuation function, we can construct a similar counterexample without relying on tie-breakers.)

By exhaustively checking all possible cases of which two D_k^i intervals contain D^i for each $i \in \{1,2\}$, one can verify that there are only two possibilities that yield districts of the correct size: either $D^1 = D_6^1 \cup D_4^1$ and $D^2 = D_8^2 \cup D_2^2$, or $D^1 = D_8^1 \cup D_2^1$ and $D^2 = D_6^2 \cup D_4^2$.

In the first case, this leaves

$$\overline{[0,1] \setminus (D^1 \cup D^2)} = \left[\frac{6}{40}, \frac{11}{40}\right] \cup \left[\frac{36}{40}, 1\right] \cup \left[\frac{26}{40}, \frac{34}{40}\right] \cup \left[\frac{15}{40}, \frac{18}{40}\right].$$

to be divided among the remaining 2 districts. Since each district can be the union of at most 2 intervals, each district must completely contain two of these intervals. This is impossible, since the four intervals have measure $\frac{5}{40}$, $\frac{4}{40}$, $\frac{8}{40}$, and $\frac{3}{40}$, no two of which sum to $\frac{10}{40}$.

Similarly, in the second case, we must divide

$$\overline{[0,1] \setminus (D^1 \cup D^2)} = \left[\frac{6}{40}, \frac{12}{40}\right] \cup \left[\frac{35}{40}, 1\right] \cup \left[\frac{26}{40}, \frac{33}{40}\right] \cup \left[\frac{16}{40}, \frac{18}{40}\right].$$

among the remaining 2 districts. Now the four intervals have measure $\frac{6}{40}$, $\frac{5}{40}$, $\frac{7}{40}$, and $\frac{2}{40}$. Again, no two these sum to $\frac{10}{40}$, so it is impossible to complete the partition in this case as well.

Thus, no GT partition exists. $\square$

EC.6. Finding X Sets in Practice

Our algorithm to find the real-world X sets from Figure 5 is as follows. For each state, we first compute the target size $|X|$ using data from our simulations in Section 4:

$$|X| := (\text{max fraction of Democratic districts}) - (\text{min fraction of Democratic districts}).$$

For example, in the state of Texas, the highest observed number of seats won by Democrats was 19 and the lowest was 11, so X should contain a population equivalent to $|X| = (19-11)/36 = 2/9$ of the total state.

Our objective is then to find a subgraph of the desired size in which the Democratic vote share is 50%. To accomplish this, we use a variant of the ReCom algorithm, which typically leads to nicely-shaped regions: We repeatedly draw a random spanning tree of the entire graph and see if there exists an edge whose removal splits off a set X of the target size with Democratic vote share in the range $[49.5\%, 50.5\%]$. We then run a biased flip chain to equalize the vote shares up to the size of a single precinct/VTD, which kept the target Democratic vote share within the target range.

We ran this algorithm six times on each of the five states: GA, NC, PA, TX, and VA. Each run drew random spanning trees until it could break off a sub-tree of the target population $|X|$ (within 1%) with the target Democratic vote share (within 0.5%). Two of the six runs targeted a Democratic vote share of 50%, two of the runs targeted 47%, and two of the runs targeted 53%. These regions are shown in Figures EC.2-EC.6. Histograms of the vote share distributions are shown in Figure EC.7.

EC.7. Empirical results omitted from Section 4.2

In Section 4.2 we report the effect of enforcing the geometric target constraint on competitiveness and compactness in Texas. In Figures EC.8-EC.10 we report the full results for the range of deviations considered.

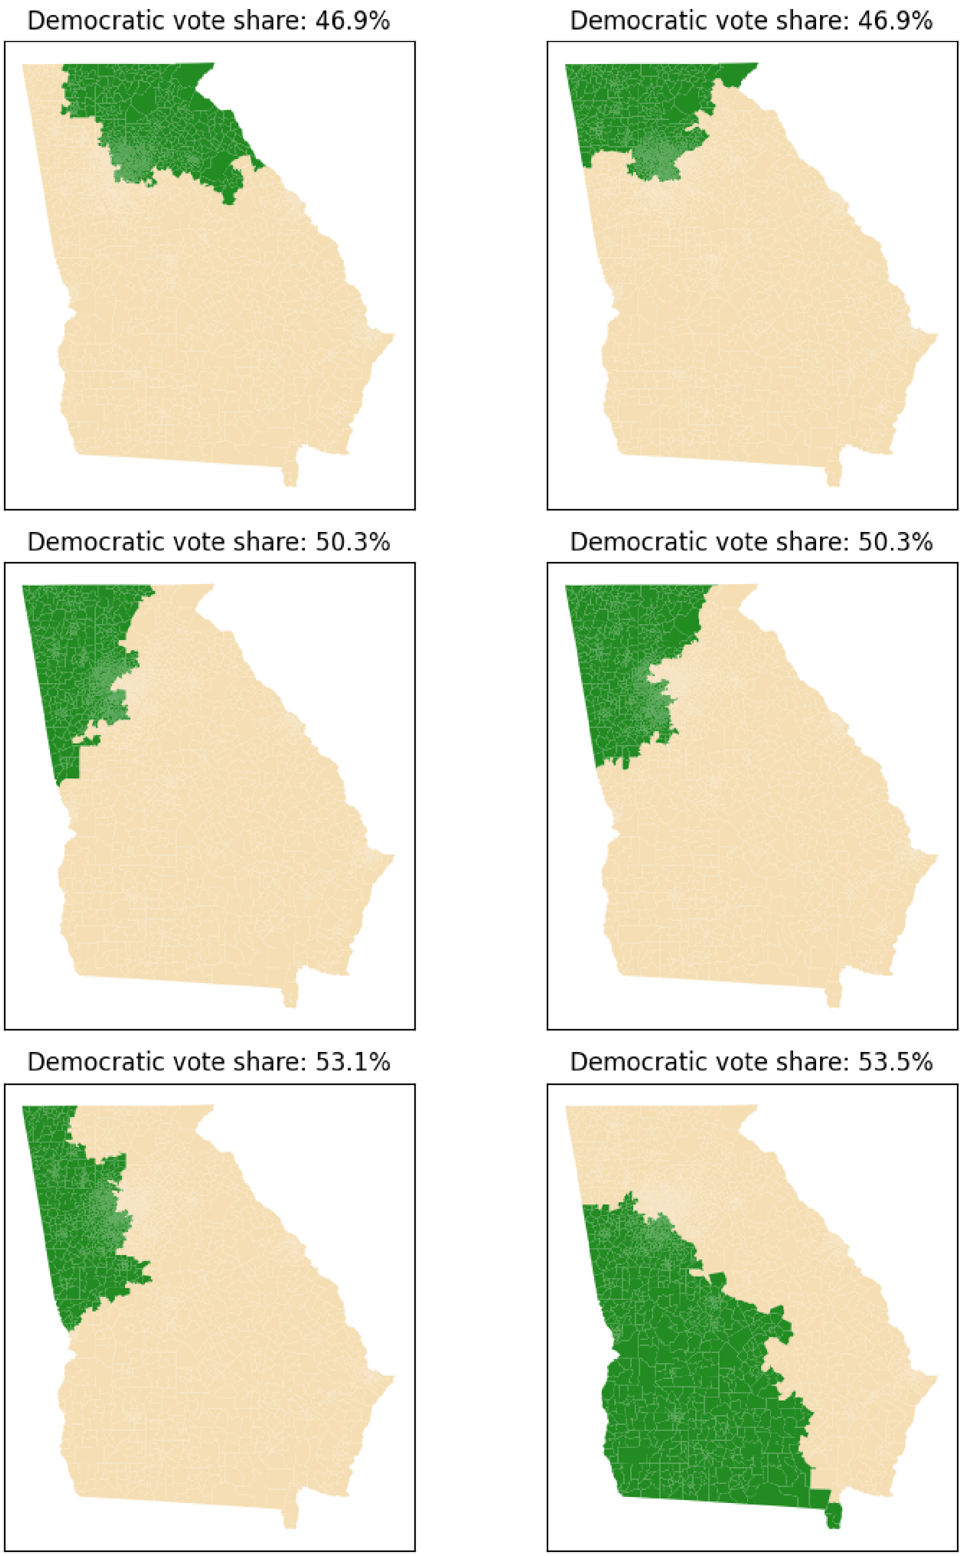

Figure EC.2 Georgia - Random regions of size $|X| = 5/14$ with various Democratic vote shares.

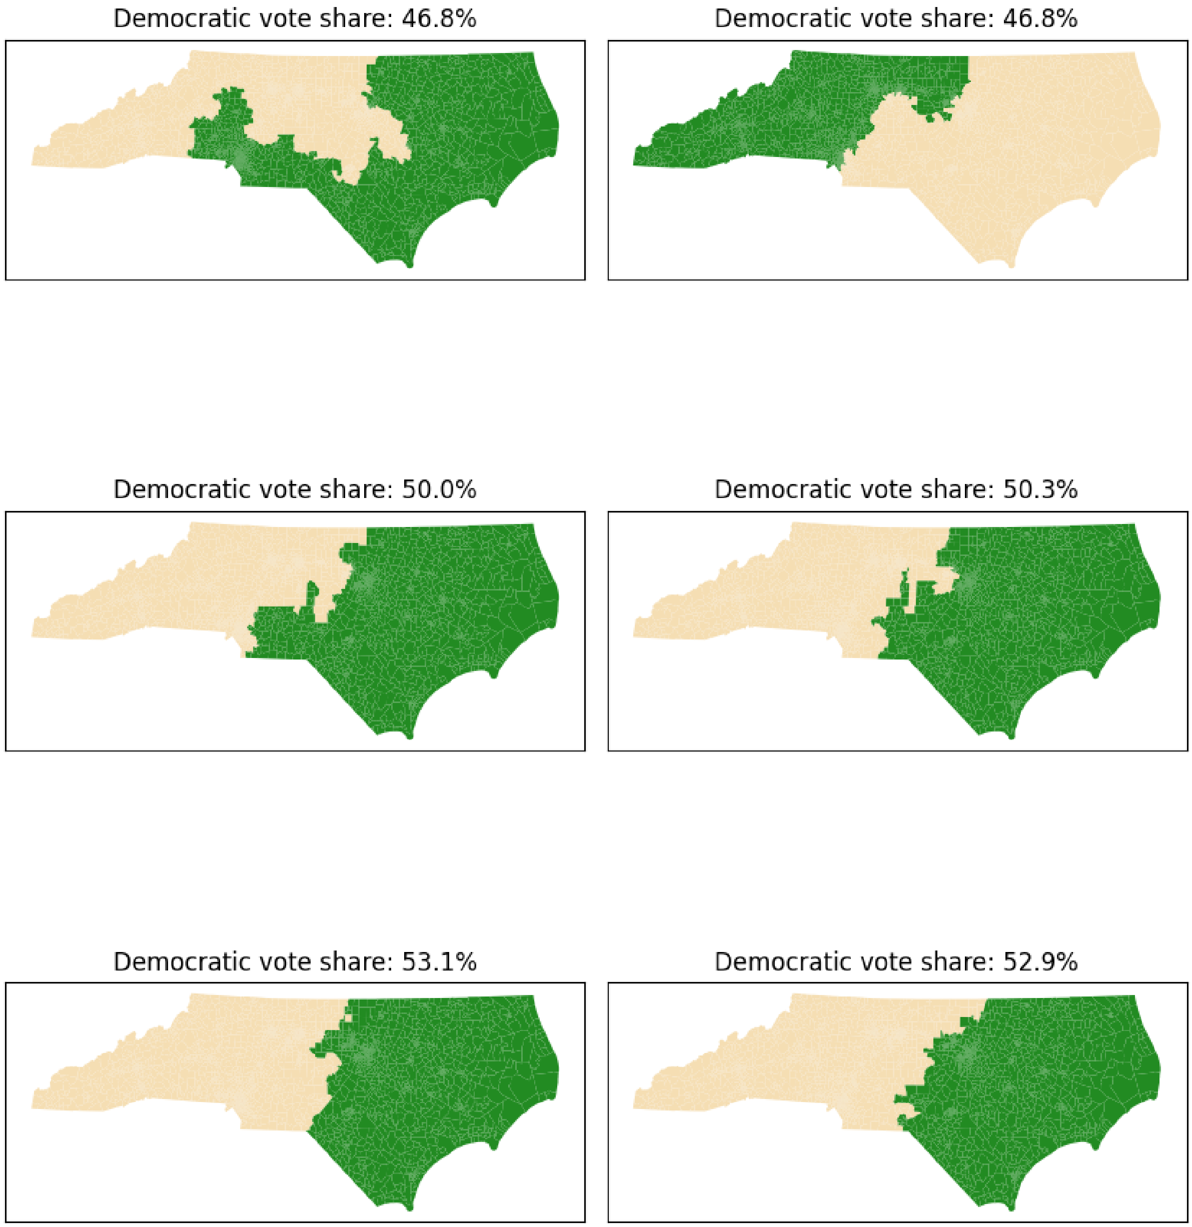

Figure EC.3 North Carolina - Random regions of size $|X| = 6/13$ with various Democratic vote shares.

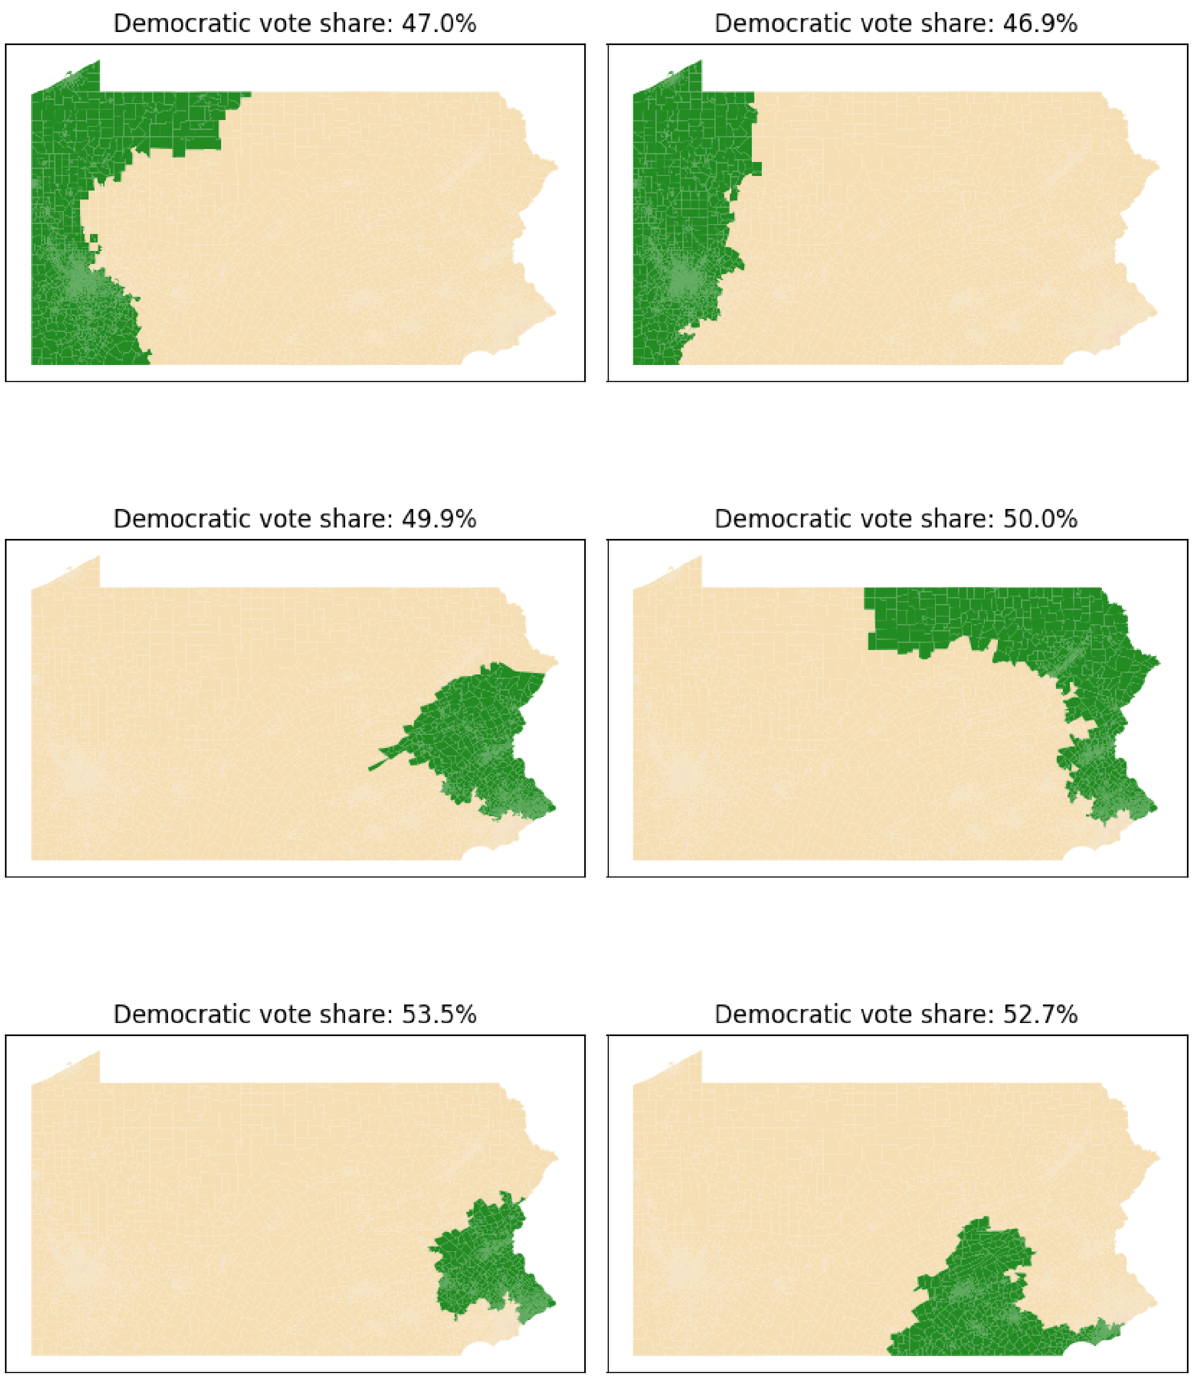

Figure EC.4 Pennsylvania - Random regions of size $|X| = 4/18$ with various Democratic vote shares.

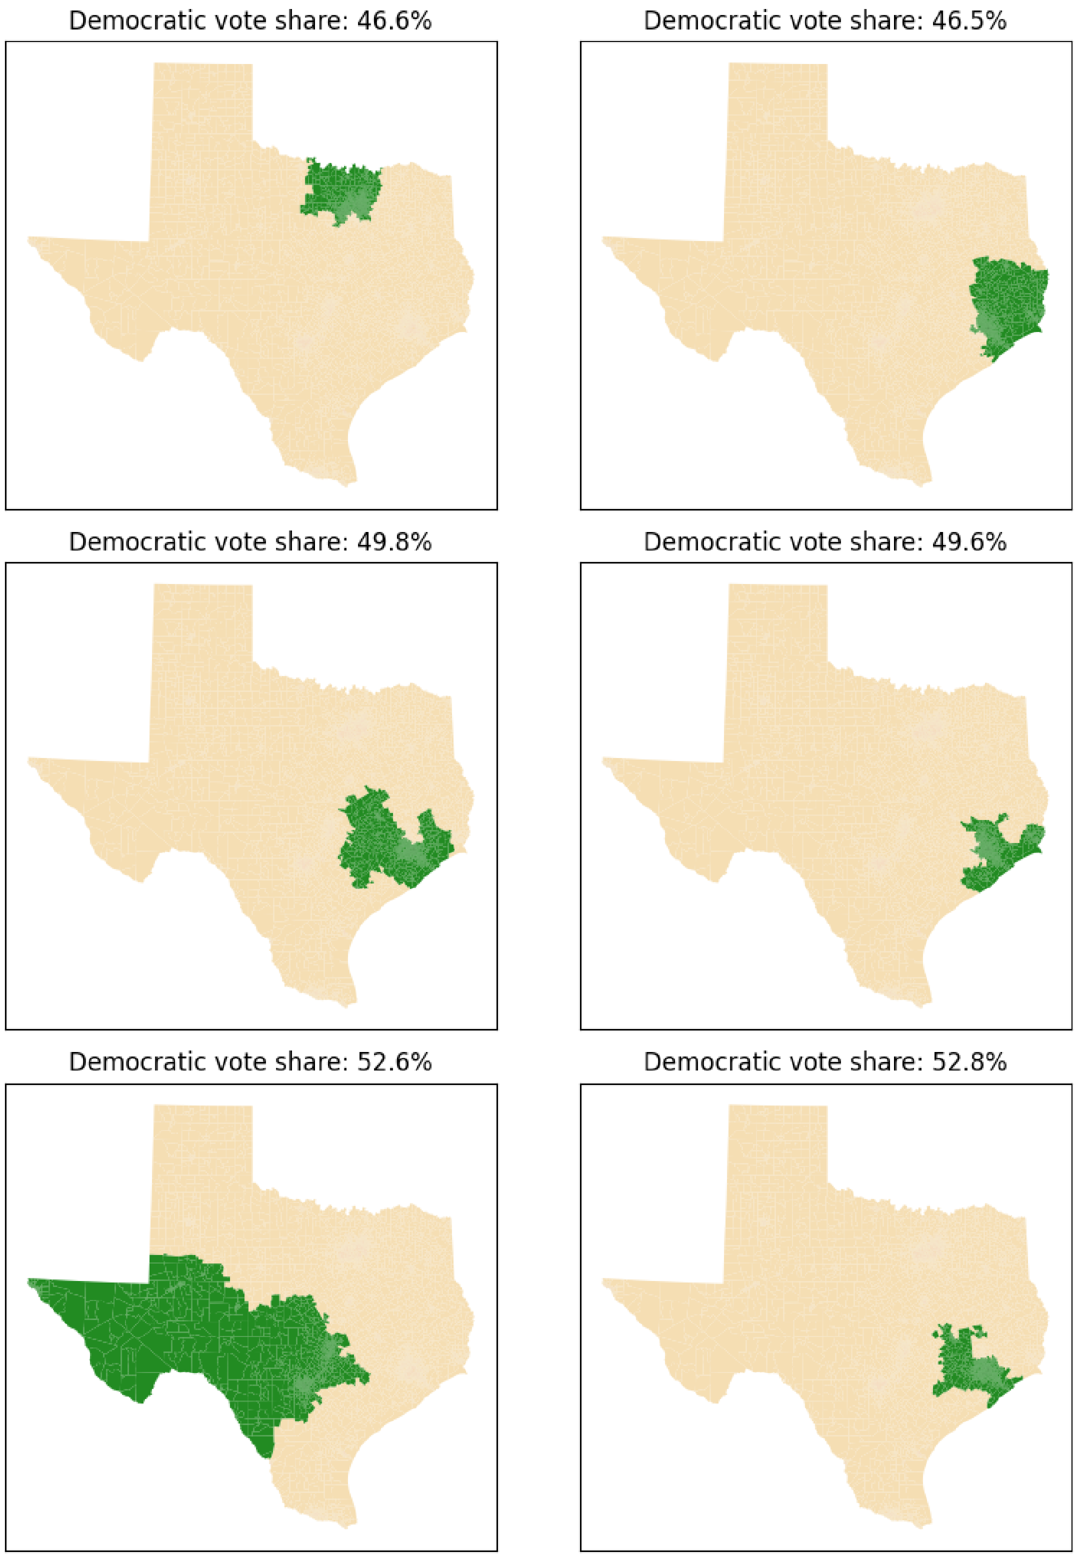

Figure EC.5 Texas - Random regions of size $|X| = 8/36$ with various Democratic vote shares.

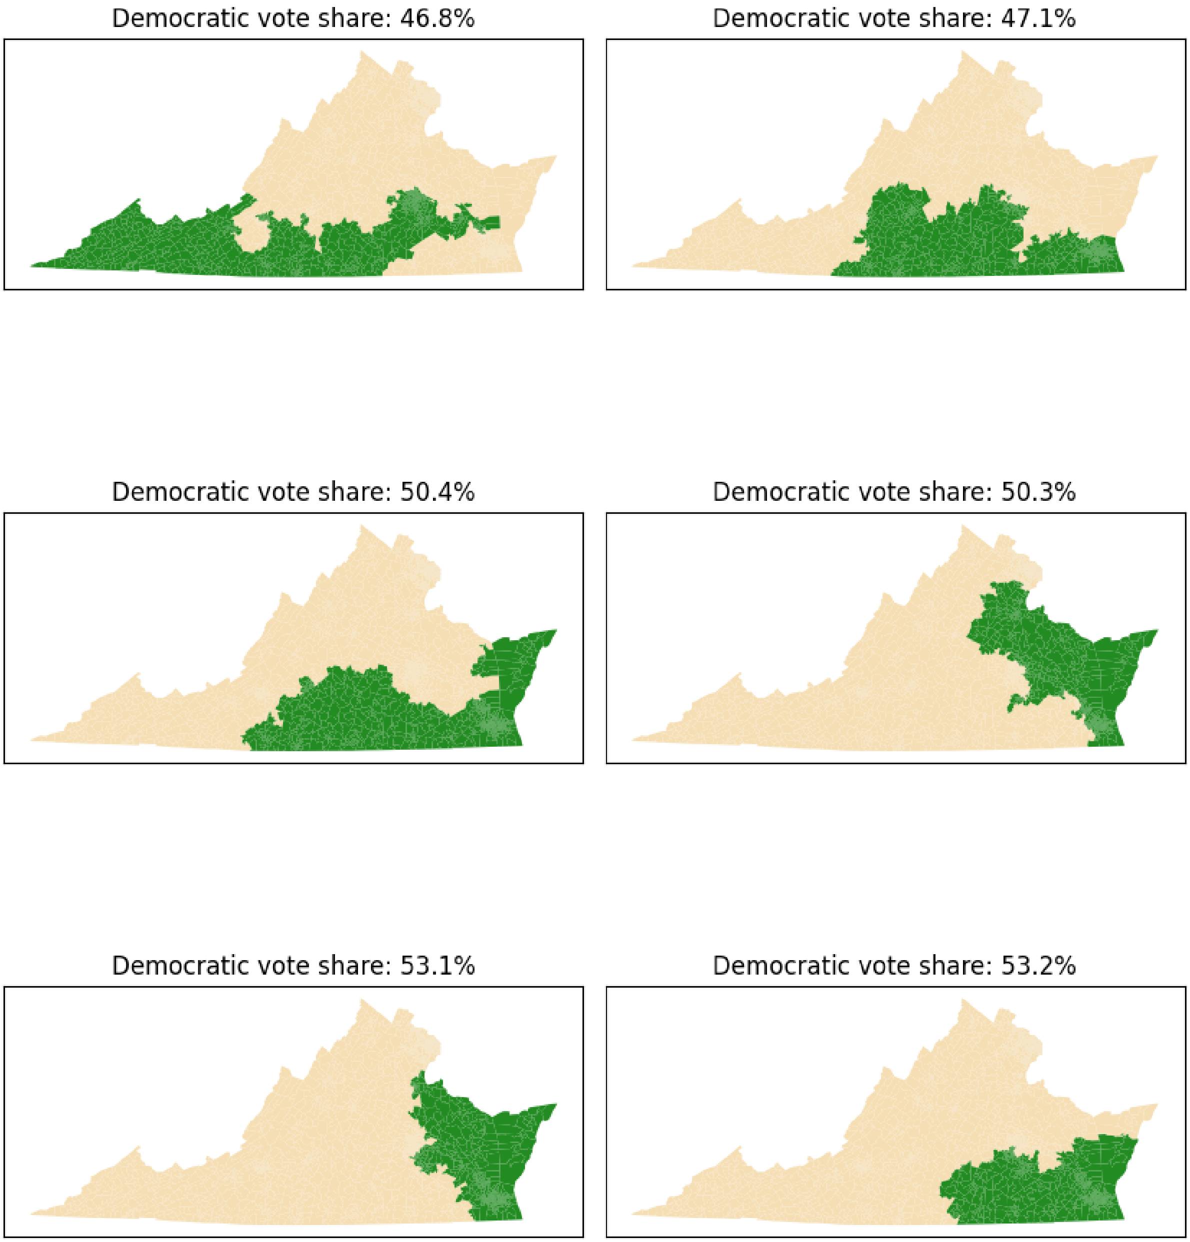

Figure EC.6 Virginia - Random regions of size $|X| = 3/11$ with various Democratic vote shares.

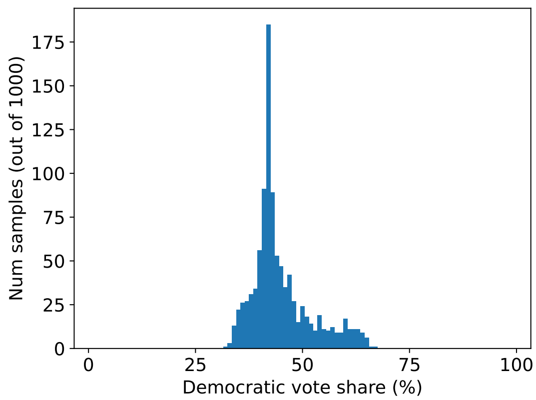

Georgia

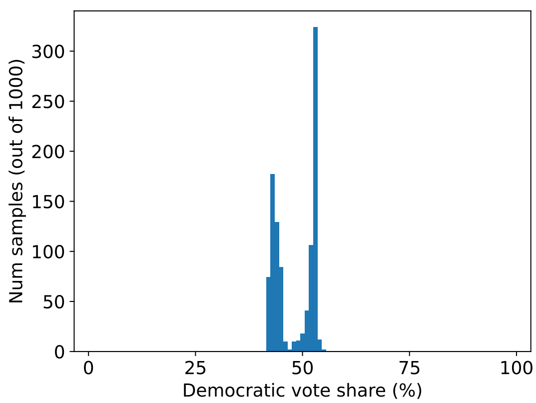

North Carolina

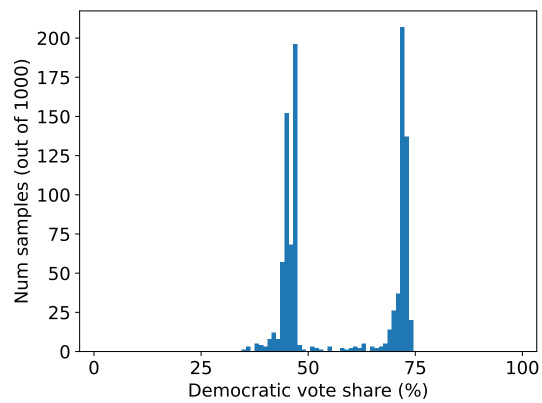

Pennsylvania

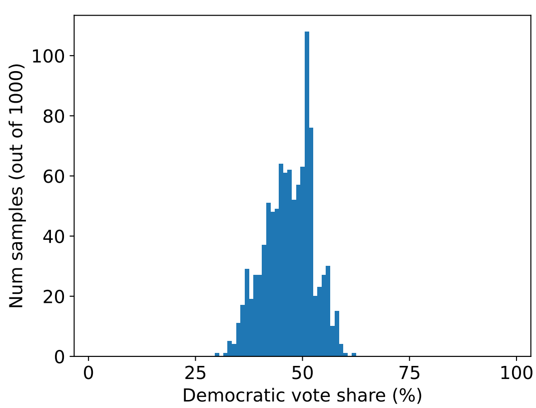

Texas

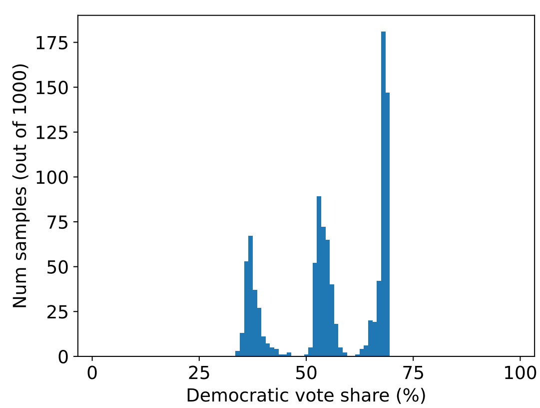

Virginia

Figure EC.7 Distributions of the Democratic vote share in random regions of the target size $|X|$.

Figure EC.8 The largest number of competitive districts among GT partitions compared to the maximum observed (black dotted line) for each of the deviations considered. The color of the bar represents which party deviates. The golden bar reports the case when neither party deviates.

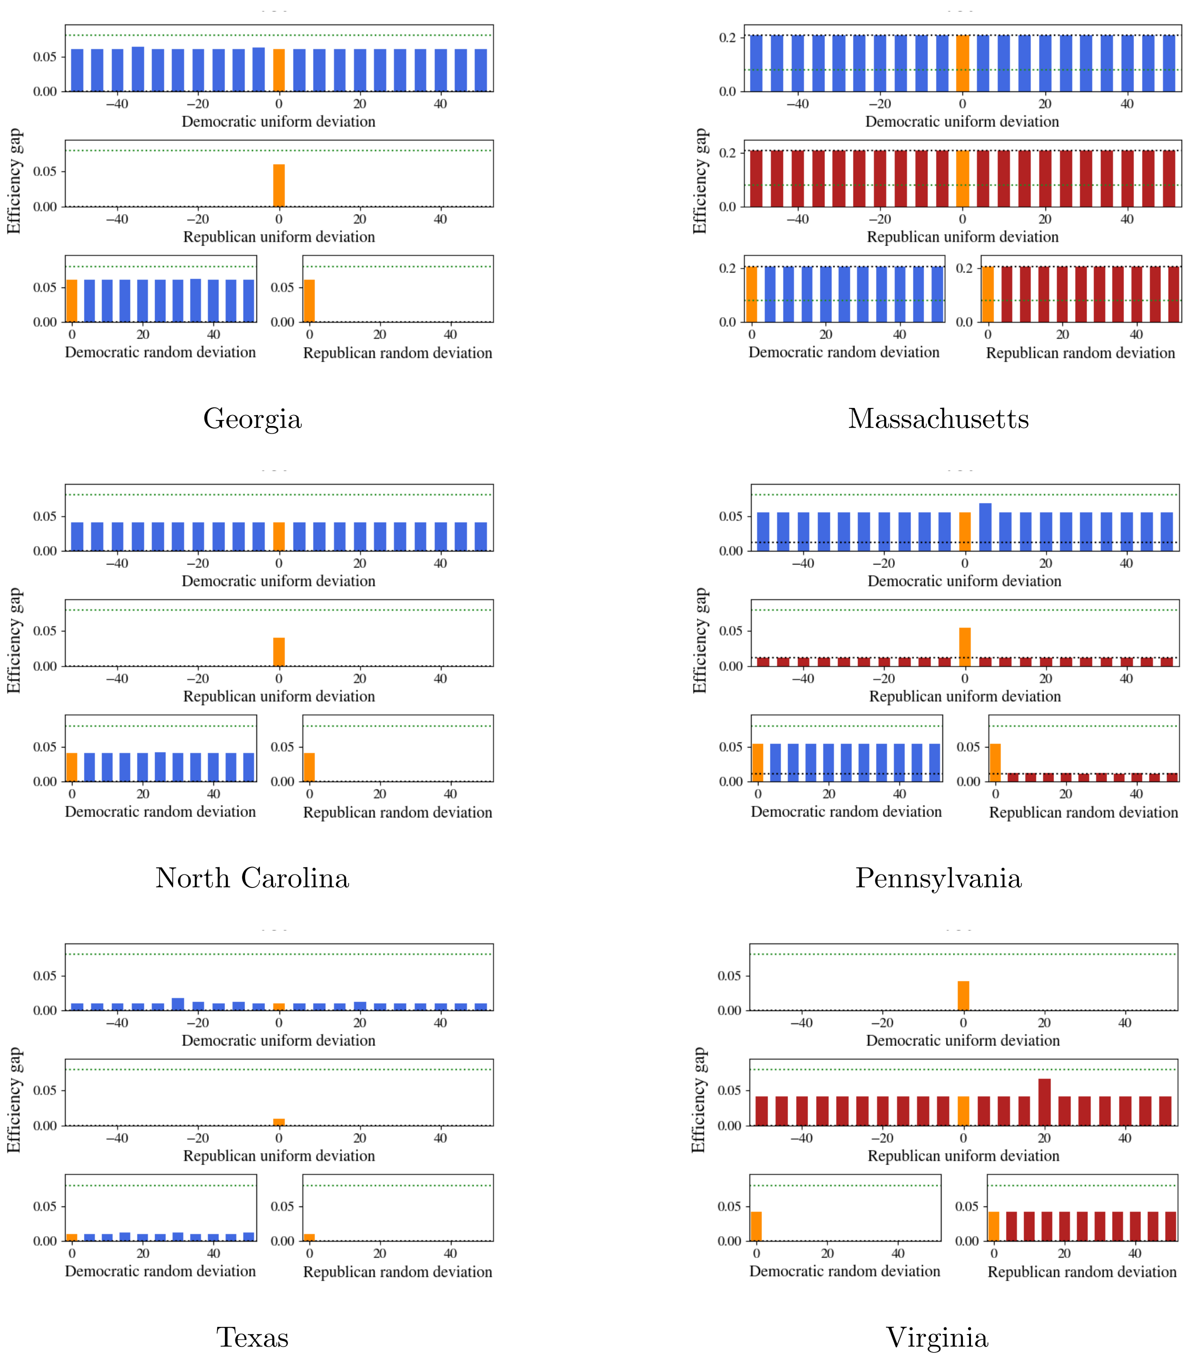

Figure EC.9 The smallest absolute efficiency gap among GT partitions compared to the best observed efficiency gap (black dotted line), and a threshold of 8% (green dotted line). The color of the bar represents which party deviates. The golden bar reports the case when neither party deviates.

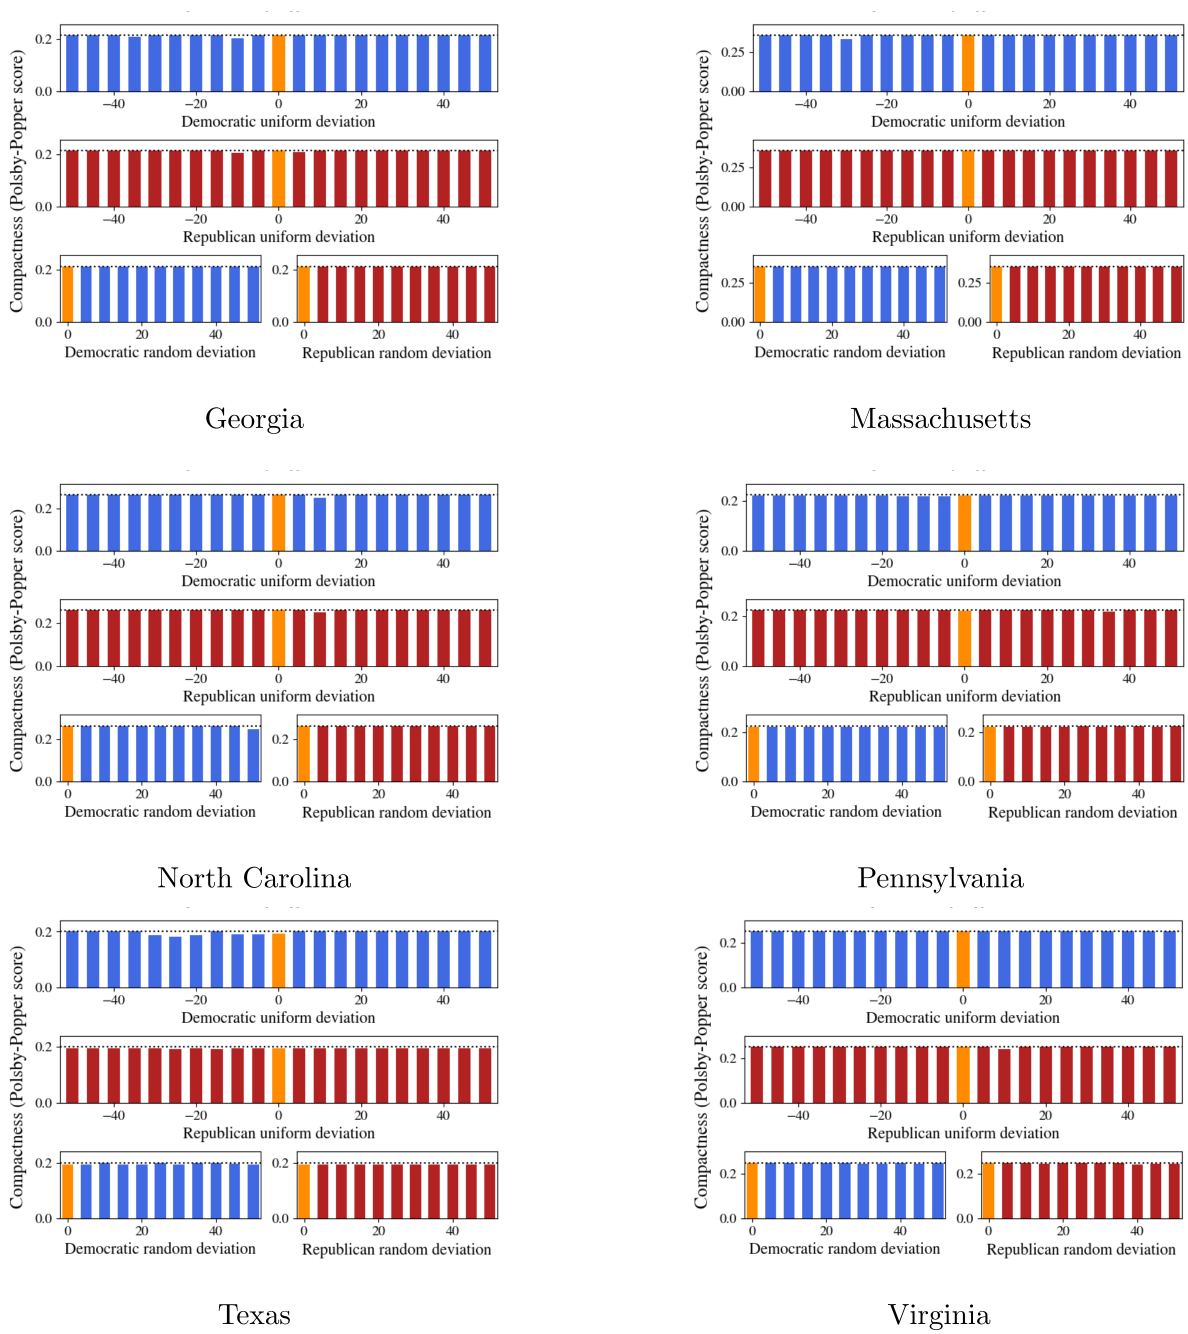

Figure EC.10 The most compact GT partitions compared to the best Polsby-Popper score observed (black dotted line) for each of the deviations studied. The color of the bar represents which party deviates. The golden bar reports the case when neither party deviates.