A widespread belief about county splits in political districting plans is wrong

Austin Buchanan, Soraya Ezazipour, Maral Shahmizad
Industrial Engineering & Management, Oklahoma State University, {buchanan,sezazip,maral.shahmizad}@okstate.edu

Consider the task of dividing a state into k contiguous political districts whose populations must not differ by more than one person, following current practice for congressional districting in the USA. A widely held belief among districting experts is that this task requires at least $k-1$ county splits. This statement has appeared in expert testimony, special master reports, and Supreme Court oral arguments. In this paper, we seek to dispel this belief. To illustrate, we find plans for several states that use *zero* county splits, i.e., all counties are kept whole, despite satisfying contiguity and 1-person deviation. This is not a rare phenomenon; states like Iowa and Montana admit hundreds, thousands, or tens of thousands of such plans. In practice, mapmakers may need to satisfy additional criteria, like compactness, minority representation, and partisan fairness, which may lead them to believe $k-1$ splits to be minimum. Again, this need not be true. To illustrate, we conduct short case studies for North Carolina (for partisan fairness) and Alabama (for minority representation). Contrary to expert testimony and Supreme Court oral arguments from *Allen v. Milligan* (2023), we find that fewer than $k-1$ county splits suffices, even when subjected to these additional criteria. This demonstrates our narrow point that $k-1$ county splits should not be assumed minimum and also suggests that districting criteria do not conflict as much as people sometimes believe. The optimization methods proposed in this paper are flexible and can assist mapmakers in satisfying them.

Key words: political districting, county splits, political subdivisions, integer programming, optimization

---

1. Introduction

The vast majority of US states require the preservation of political subdivisions (e.g., counties, cities, towns) in their political districts; this is true for both congressional and legislative districts (NCSL 2021). Arguably, the most popular way to quantify splitting is the number of splits (Carter et al. 2020, Cervas and Grofman 2020, Autry et al. 2021, Nagle 2022, Dave's Redistricting App 2024, Shahmizad and Buchanan 2023), which is (nearly) equivalent to the number of parts or pieces (Gladkova et al. 2019, Becker and Gold 2022), intersections (Wachspress and Adler 2021), or traversals (Carter et al. 2020). For example, if a county is wholly assigned to one district, then it contributes zero county splits. If it is divided across two districts, then it contributes one split. Generally, if a county is divided across k districts, then it contributes $k-1$ county splits. Usually, the sum total number of county splits is reported.

The academic literature on redistricting makes several claims about the number of county splits s and how this quantity relates to the number of districts k. Often, the claim is that, in any districting plan, the number of splits is at least the number of districts minus one, i.e., $s \geq k-1$, especially if districts must not differ in population by more than one person (Autry et al. 2021, Nagle 2022), which is the norm for congressional districting (NCSL 2020, 2023). Sometimes, it is further asserted that the *minimum* number of splits s^* precisely achieves this quantity for (almost) all instances, i.e., $s^* = k-1$, see Nagle (2022)[1].

These claims have been repeated in court cases by a wide variety of districting experts, including in the Supreme Court case *Allen v. Milligan* (2023). In it, Alabama's congressional districts were challenged under Section 2 of the Voting Rights Act (VRA) for diluting the voting strength of Black voters. Below, we provide excerpts from expert testimony, cross examination, Supreme Court oral arguments, and the Special Master's report. These quotes show that many people involved in the case (from all sides) believe that drawing seven districts requires six county splits.

- From expert testimony (*Allen v. Milligan* 2021):

In order to make seven finely population-tuned districts, it is necessary to split at least six of Alabama's 67 counties into two pieces, or to split some counties into more than two pieces.

- From cross examination before a three-judge district court (*Allen v. Milligan* 2022a):

Q: *At least six times, a county must be split to get the one person one vote minimal deviation that we're looking for, right?*

A: *I think a precise way to phrase it would be that there have to be at least six additional county pieces as a way of phrasing.*

Q: *And that's simple math that counties rarely line up where– you're unlikely to have a county that's exactly 717,000 whatever people in it to form that one perfect district, so you are going to probably have to split it at least a little to equalize it, right?*

A: *That's the idea, yes.*

- During Supreme Court oral arguments (*Allen v. Milligan* 2022b):

JUSTICE KAVANAUGH: *. . . you look at respecting county lines, for example, right? That's an important one. And this did. This new district did just as well, if not better, in respecting county lines. At least that's the argument. So I want to hear your response to that. . .*

[1] Nagle's work is unpublished but has nevertheless been impactful. For example, the Analyze tab on Dave's Redistricting App (2024) states that "Given k districts, you might need to split counties $k-1$ times for district populations to be 'roughly' equal". A developer of DRA attributes this claim to Nagle and reiterated to us that $k-1$ is *minimum* (Ramsay 2022). In another example, Nagle's work was favorably referenced by the special master Cervas (2022) in *Harkenrider v. Hochul* (2022), who redrew New York's congressional and state senate districts after they were found to be unconstitutional Democratic partisan gerrymanders.

MR. LACOUR: *Well, three of the Duchin plans split more counties than necessary. The Cooper plans keep them together but the same number of splits. Six is the minimum you have to have.*

- From the Special Master's report, whose remedial plans all have at least six county splits (*Allen v. Milligan* 2023):

Second, to minimize county splits, the Special Master proposes placing Elmore County... entirely in District 6... Finally, after avoiding county splits where possible, the Special Master also sought to minimize the number of split precincts...

In this paper, our aim is to dispel these beliefs. We make three main points:

1. Often, fewer than $k-1$ county splits suffice to satisfy the most basic districting criteria (i.e., 1-person deviation and contiguity). For example, we show that several states (Idaho, Iowa, Mississippi, Montana, Nebraska, West Virginia) can do so using *zero* county splits.

2. These examples are not rare flukes. For example, Montana admits 30,223 contiguous, whole-county plans with 1-person deviation, and Iowa admits more than 1,000 such plans.

3. Even when constrained by other criteria (e.g., compactness, minority representation, partisan fairness), $k-1$ need not be the minimum number of county splits. For example, we provide a reasonably configured plan for Alabama with two majority-Black districts and 1-person deviation that nevertheless exhibits fewer than $k-1$ county splits. Similarly, we provide a reasonably configured plan for North Carolina that scores well on partisan fairness metrics, despite satisfying 1-person-deviation and exhibiting fewer than $k-1$ county splits.

We conclude that $k-1$ county splits should not be assumed minimum. Going forward, districting experts should either remain agnostic to such statements, or rigorously prove or disprove them using exact methods like ours. Our case studies also suggest that districting criteria do not conflict as much as people sometimes believe. The optimization methods proposed in this paper are inherently flexible and can assist mapmakers in satisfying them. To this end, our Python codes have been publicly released under the GPL-3.0 license, allowing anyone to run, study, share, or modify them. A more detailed accounting of our contributions is given in Section 2.4, but first we provide important background and context for these contributions.

2. Background and Literature Review

Districting problems are usually cast in terms of graphs. To wit, let $G=(V,E)$ be a graph whose vertices V represent a state's geographic units, which could be counties, census tracts, voting precincts, census blocks, etc. The edges E indicate which pairs of geographic units are adjacent on the map.

We seek to partition the state into k districts $(D_1, D_2, \ldots, D_k)$. Alternatively, we can think of a districting plan as a function $d: V \to [k]$ that maps each vertex to a district number from the set $[k] := \{1, 2, \ldots, k\}$. We use both representations interchangeably.

Usually, each district is required to be contiguous on the map; in graph terms, this means that each district $D \subseteq V$ should induce a subgraph $G[D] = (D, E \cap \binom{D}{2})$ that is connected, where $\binom{D}{2}$ denotes the collection of two-element subsets of D. Each geographic unit $i \in V$ has an associated population p_i. When $S \subseteq V$ is a subset of vertices, we use $p(S)$ as a shorthand for its population $\sum_{i \in S} p_i$. Each district population should be near to the ideal population $p(V)/k$, say, at least L and at most U. In 1-person-deviation, these population bounds are $L = \lfloor p(V)/k \rfloor$ and $U = \lceil p(V)/k \rceil$, where $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ are the floor and ceiling functions, respectively. Setting L and U to these values ensures that the *overall population deviation* or *total population deviation*, which is the largest district population minus the smallest district population, i.e.,

$$\max\{p(D_j) \mid j \in [k]\} - \min\{p(D_j) \mid j \in [k]\},$$

will be at most one. This is the usual way to quantify deviation in court cases (Hebert et al. 2010).

2.1. Norms Around Population Balance

Up until the 1960s, political districts in the USA sometimes had highly disparate populations. For example, after the 1960 census, Georgia's congressional districts varied in population from 272,154 at the low end to 823,680 at the high end. As a result, voters in the least-populous district had three times the voting strength as those in the most-populous district. In *Wesberry v. Sanders* (1964), the US Supreme Court overturned these congressional districts for violating Article 1, Section 2 of the US Constitution. This decision was part of a broader "one-person, one-vote" revolution that included other landmark cases such as *Baker v. Carr* (1962) and *Reynolds v. Sims* (1964), which concerned state legislative districts in Tennessee and Alabama, respectively, and whose decisions were instead based on the Equal Protection Clause of the 14th Amendment.

In the years since, different norms have evolved for congressional districts and state legislative districts, due to their different legal footings (i.e., Article 1, Section 2 versus Equal Protection), with the standards for congressional districts being considerably stricter. Nowadays, most states enact congressional plans with 1-person deviation. For example, after the 2010 and 2020 censuses, $29/43 \approx 67\%$ and $24/44 \approx 54\%$ of states did so, respectively (NCSL 2020, 2023), if we exclude states with just one congressional district. Some may see this insistence on 1-person deviation as silly given that: (1) census undercounts and overcounts are orders of magnitude larger (Wang 2022a,b), and that (2) district populations change considerably between censuses, a phenomenon that can lead to *within-cycle malapportionment* (DeFord et al. 2023). Indeed, the data is already more than two years old when the new districts are first used for elections, and these districts are typically used for ten years thereafter. However, there is a reason states often follow this practice. Any deviation from precise mathematical equality risks a lawsuit and must be justified. As the Supreme Court held in *Karcher v. Daggett* (1983):

"Parties challenging apportionment legislation bear the burden of proving that population differences among districts could have been reduced or eliminated by a good faith effort to draw districts of equal population. If the plaintiffs carry their burden, the State must then bear the burden of proving that each significant variance between districts was necessary to achieve some legitimate goal."

In the years following *Karcher*, congressional plans with tiny deviations have been rejected by federal courts, including a Pennsylvania plan with a 19-person deviation (i.e., 0.0029%) (Hebert et al. 2010). However, larger deviations approaching 1% have sometimes been permitted with sufficient justification. For example, West Virginia's congressional districts were upheld in *Tennant v. Jefferson County* (2012) despite a 4781-person deviation, justified by the state's desire to keep counties whole. However, most states opt for 1-person deviation to be safe.

Despite the emphasis on 1-person deviation in practice, nearly all computational redistricting work imposes larger deviations, even when drawing congressional districts. For example, in the recent optimization literature, Swamy et al. (2023) use either a ±2%, ±1%, or ±0.25% deviation, Dobbs et al. (2023a) use a ±1% deviation, Validi et al. (2022), Validi and Buchanan (2022), Zhang et al. (2024) and Dobbs et al. (2023b) use a ±0.5% deviation, and Shahmizad and Buchanan (2023) primarily use a ±0.5% deviation. Ensemble methods, which aim to understand the underlying *distribution* of plans, also use larger deviations: DeFord and Duchin (2019) primarily use a ±2% deviation (but sometimes tighten or relax this to ±0.5% or ±5%); Becker et al. (2021) use a ±1% deviation; DeFord et al. (2021) appear to use ±5% deviation (see Fig. 9); Autry et al. (2021) use a ±2% deviation; McCartan et al. (2022) and Kenny et al. (2023) primarily use a ±0.5% deviation; McCartan and Imai (2023) use a ±0.1% deviation and remark that using a tighter deviation would require them to use more granular geographic units like census blocks rather than precincts. Usually, these larger deviations are justified by statements like that of DeFord et al. (2021):

"Even for Congressional districts, which are often balanced to near-perfect equality in enacted plans, a precinct-based ensemble with $\leq 1\%$ deviation can still provide a good comparator, because those plans typically can be quickly tuned by a mapmaker at the block level without breaking their other measurable features."

A notable exception is the computational work of Cohen-Addad et al. (2018) that *does* achieve 1-person deviation, but does so with block-level tuning. Census blocks often correspond to city blocks in urban areas and are considerably more granular than the other geographic units (e.g., counties, tracts, precincts) usually used in computational redistricting works. For example, Alabama has 67 counties, 1437 tracts, 1837 precincts, and 185,976 blocks. As suggested by the DeFord et al. quote above, mapmakers may first draw rough plans using larger geographic units and subsequently break them into census blocks for final tuning (e.g., to achieve 1-person deviation). As one might expect, this final step may cause additional splits to political subdivisions like counties. So, the emphasis on 1-person deviation and county preservation is a distinguishing feature of the present work.

2.2. County Splitting Scores

Each state is subdivided into a set of counties C or county equivalents. Similarly, each county $c \in C$ is subdivided into precincts, tracts, or blocks, the set of which will be denoted by V_c, with $V_c = \{c\}$ if G is itself a county-level graph.

Let $d : V \to [k]$ be a districting plan. If $S \subseteq V$ is a subset of vertices, then the set of districts that S is assigned to is denoted by $d[S] := \{d(i) \,:\, i \in S\}$; this is simply the image of set S under d. In particular, county c's vertices V_c are assigned to the districts $d[V_c]$.

Definition 1 (county splits). The county c is *whole*, *intact*, or *preserved* in plan d if $d[V_c]$ is a singleton (i.e., $|d[V_c]| = 1$), in which case it contributes zero splits; otherwise, it is *split*, with the number of splits being $|d[V_c]| - 1$. The (total) number of county splits is $\sum_{c \in C}(|d[V_c]| - 1)$.

The number of county splits should not be confused with the *number of split counties* (Wachspress and Adler 2021) or *number of counties split* (Becker and Gold 2022), which are the size of the set $C_{\text{split}} = \{c \in C : |d[V_c]| > 1\}$. Other splitting scores include the number of *parts* (Gladkova et al. 2019) or *intersections* (Wachspress and Adler 2021), which are $\sum_{c \in C} |d[V_c]|$, i.e., the number of county splits plus the constant $|C|$. Thus, the differences between county splits, parts, and intersections are only cosmetic, and they are all equivalent from an optimization perspective. More complicated splitting scores exist, including various entropy-based scores (Becker and Gold 2022, Guth et al. 2022) and the number of *pieces* (Gladkova et al. 2019) or *fragments* (Becker and Gold 2022), which count the number of connected components of the intersections $V_c \cap D_j$ between each county c and each district D_j; McCartan and Imai (2023) use this same score minus the number of counties and call it *splits*. For more, we refer the reader to Becker and Gold (2022).

2.3. Claims about County Splits

Below we review some common claims about the minimum number of county splits s^* and how this quantity relates to the number of districts k and to the maximum number of county clusters c^* (defined later). In short, redistricting folklore states that $s^* \geq k - 1$ or $s^* \leq k - 1$ or possibly both (i.e., $s^* = k - 1$), none of which are generally correct. Meanwhile, researchers like Carter et al. (2020) state that $s^* = k - c^*$, which has been confirmed to hold in practice by Shahmizad and Buchanan (2023) for congressional districts ($\pm 0.5\%$ deviation) and legislative districts ($\pm 5\%$ deviation), but is also generally incorrect. However, it is indeed always true that $s^* \geq k - c^*$.

$k - 1$ is neither an upper nor lower bound. Redistricting folklore states that, when dividing a state into k contiguous and population-balanced districts, $k - 1$ county splits suffice. For some intuition, consider four counties arranged in a line, each with a population of 75, as in Figure 1.

Suppose we seek $k = 3$ equipopulous districts. We may create our first district with the leftmost county (population 75) and add to it 25 people from the second county, introducing one split. Then, create our second district with the remaining 50 people from the second county and add to it 50 people from the third county, introducing a second split. Then, create the third district from the remaining 25 people and the entire rightmost county. Thus, we have created three districts using two county splits, as folklore would suggest. Of course, this idea applies to more complicated instances; the important assumption is that we should be able to carve $k-1$ districts from the state, one-by-one, each time introducing one county split, and take what remains as the final district.

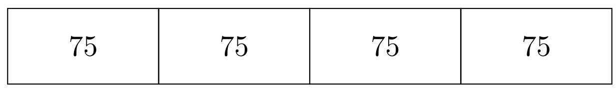

Figure 1 A hypothetical districting instance with four counties in a line

This is sometimes impossible, as noted by Carter et al. (2020). Here, we give a modified example from Shahmizad and Buchanan (2023). Consider a hub county with 35 people that is adjacent to three spoke counties, each with a population of 55, as in Figure 2. Suppose we are to divide this state into two districts, each with a population between 95 and 105. We must split at least one of the spoke counties (otherwise, all spoke counties will be kept whole and some district will contain at least two of them, causing its population to reach 110, which is too much). Now, each district can take at most 55 people from this split spoke county, which is too little, meaning that each district must extend into the hub county, splitting it as well. Thus, we need at least two county splits, which is more than $k-1$. Shahmizad and Buchanan extend this example to show that any number of splits $k+q$ might be needed, for any nonnegative integer q. So, there is generally no way to upper bound the minimum number of splits s^* by a function of the number of districts k.

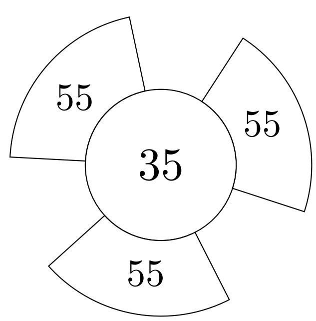

Figure 2 A hypothetical districting instance with a hub county and three spoke counties

Further, the academic literature often claims $k-1$ to be a lower bound. Nagle (2022) states that forcing districts to satisfy a 1-person deviation makes it “highly probable that the minimum number of county splits is uniquely given as the number of districts minus one". Likewise, Autry et al. (2021) consider this a "reasonable" assumption.

$k - c^*$ is usually minimum in practice (but not always). Carter et al. (2020) propose a more nuanced claim. They recognize that the number of county splits is sometimes less than $k-1$.

To illustrate, consider dividing Alabama's total population of 5,024,279 across seven districts, so each has an ideal population of 5,024,279/7≈717,754.14. Thus, to achieve a 1-person deviation, there must be six districts with a population of L=717,754 and one district with population U=717,755. It turns out that Alabama's counties can be partitioned into two contiguous sets, one with a population of $L+U$ and another with a population of $5L$, see Figure 3. We can consider them as two separate, miniature districting instances, the first with two districts and the second with five districts. We will see that we can divide up the first using one county split and the second using four county splits, for a total of five county splits. So, by first dividing the state's counties into two miniature districting instances, we save one county split (beyond the folklore $k-1$ number).

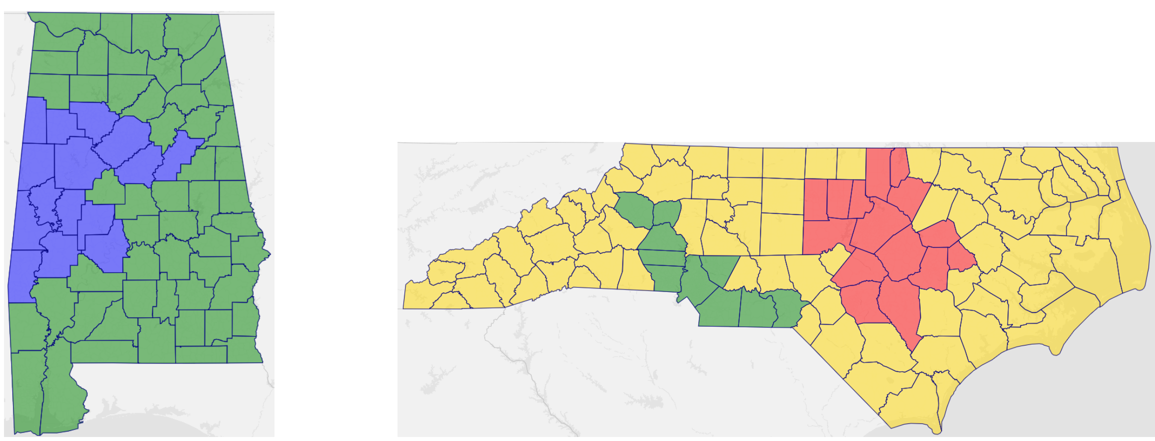

Figure 3 County clusterings for Alabama and North Carolina with two and three clusters, respectively

In another example, consider dividing North Carolina's total population of 10,439,388 across fourteen districts, so each has an ideal population of roughly 745,670.57. Thus, to achieve a 1-person deviation, there must be six districts with a population of L =745,670 and eight districts with population U =745,671. It turns out that North Carolina's counties can be partitioned into three contiguous sets, the first with a population of $3U$, the second with population $3L+U$, and the third with population $3L+4U$. We can consider them as three separate districting instances, with three, four, and seven districts, respectively. Later, we will divide them up using two, three, and six county splits, respectively, giving eleven total splits, *two* less than the folklore $k-1$ number.

More generally, the idea behind Carter et al.'s claim is that, by first partitioning a state's counties into a maximum number c^* of miniature districting instances, we can save c^*-1 county splits, thus giving $(k-1)-(c^*-1)=k-c^*$ county splits. Indeed, in their "basic" theorem, they propose the bold claim that the minimum number of county splits s^* equals $k-c^*$. In a subsequent "enlarged" theorem, they add that the caveat that this equality holds "except in rare circumstances". Their theorem statement does not specify *what* these rare circumstances are, nor do they establish *how rare* they are in practice.

Before continuing, let us formalize the idea of a "decomposition into miniature districting instances" via *county clusterings* (Carter et al. 2020), following Shahmizad and Buchanan (2023).

Definition 2 (county clustering). A *county clustering* $(C_1, C_2, \ldots, C_q)$ is a partition of the counties along with associated *cluster sizes* $(k_1, k_2, \ldots, k_q)$ such that

1. the cluster sizes are positive integers that sum to k,
2. each cluster C_j induces a connected subgraph, and
3. each cluster C_j has a population satisfying $Lk_j \le p(C_j) \le Uk_j$.

A county clustering is *maximum* if its cardinality q is largest among all county clusterings.

Shahmizad and Buchanan (2023) point out that half of Carter et al.'s theorem always holds, that is $s^* \ge k-c^*$, a result that they name *weak split duality.* Using integer programming techniques, Shahmizad and Buchanan compute a maximum number of county clusters for each congressional and legislative districting instance across the USA, thus establishing their c^* values. Then, using the inequality $s^* \ge k-c^*$, they establish a lower bound on s^*. With other integer programming techniques, they find districting plans that achieve this lower bound, thus proving optimality in terms of minimum county splits. So, we may empirically conclude that Carter et al. are right; their $s^*=k-c^*$ "theorem" does hold in practice. (Note that the hub-and-spoke instance from earlier provides a synthetic counterexample, as it has a maximum of one county cluster $c^*=1$ but requires at least $s^* \ge 2$ splits, thus giving an example where $s^* > k-c^*$.)

Shahmizad and Buchanan primarily used a $\pm 0.5\%$ deviation for congressional instances and a $\pm 5\%$ deviation for legislative instances. However, recognizing that 1-person deviation is the norm for congressional districting, they also performed a limited set of experiments for 1-person deviation, establishing that 79% of these districting instances admit a nontrivial county clustering (i.e., with $c^* \ge 2$). So, contrary to speculations by Autry et al. (2021) and Nagle (2022), it is the norm rather than a rare exception for a state to admit a nontrivial county clustering. If these county clusterings can be extended into districting plans (as folklore would suggest), then this would yield plans with 1-person deviation and fewer than $k-1$ county splits.

2.4. Our Contributions

In this paper, we go further. We show that *zero* county splits suffice for states like Idaho, Iowa, Mississippi, Montana, Nebraska, and West Virginia. Further, we show that these counterexamples are plentiful; states like Iowa and Montana admit *hundreds, thousands, or tens of thousands* of contiguous, whole-county plans with 1-person deviation. Moreover, we show that, even when plans must satisfy other criteria like compactness, minority representation, and partisan fairness, it can still be possible for states to draw plans with fewer than $k-1$ county splits, contrary to statements made by a variety of districting experts. These contributions to districting practice are enabled by our advances in mixed-integer programming (MIP) methodology, as previewed below.

Our first methodological contribution is an approach for enumerating the top t most compact county clusters. Figuring out how to do this well took real effort, as standard approaches are inapplicable or duplicate considerable effort. For example, Gurobi's approach to finding the t best solutions (via the `PoolSearchMode` parameter) is ill-equipped to handle extended formulations, a shortcoming that has been acknowledged by Gurobi's Tobias Achterberg (Buchanan 2021). Meanwhile, extended formulations are nearly ubiquitous in districting models, including all that use the popular flow-based contiguity constraints of Shirabe (2005, 2009). We have found the top-t enumeration idea to be crucial when seeking to satisfy other criteria, e.g., when generating the Alabama plan in Section 3.1. Our approach to top-t enumeration is applicable beyond districting.

Our second methodological contribution is an extension to find whole-county districting *plans* with 1-person deviation. With it, we find more than 1,000 such plans for Iowa. It may not be apparent to those who have not worked on this instance, but it is a *huge* computational challenge. Optimization and districting enthusiasts had tried for several years, using many different techniques, to find a *single* Iowa plan with 1-person deviation. The fact that we can find more than 1,000 plans is a considerable computational feat. For some context, a plan with 5-person deviation, submitted by Harvard's Cory McCartan, won Dave Wasserman's redistricting challenge (Burger 2021). In this challenge, the task was to find a contiguous, whole-county plan for Iowa with minimum deviation. With our new methodology, we can find what would have been more than 1,000 winning (and optimal) entries. Our approach is a district-carving procedure, similar to that of McCartan and Imai (2023), except that it is optimization-driven and uses top-t enumeration.

3. Enumerating Top County Clusters with Integer Programming

Here we propose integer programming techniques to identify county clusters rooted at a given county. To obtain reasonably configured districts, we seek clusters that are compact in shape. Compactness can be measured by the number of cut edges emanating from the cluster (Duchin 2022), the cluster's boundary length, its Polsby-Popper score (Polsby and Popper 1991), or in many other ways (Young 1988, Niemi et al. 1990, Kaufman et al. 2021). For simplicity, we present only the cut edges model here; extending the model to capture the boundary length or Polsby-Popper score is straightforward using ideas from Validi and Buchanan (2022), Belotti et al. (2023) and is also implemented in our code, see also Buchanan (2023a). The most compact cluster (however that is measured) may be undesirable for any number of reasons. So, for sake of flexibility, our approach enumerates the top t most compact clusters, and the user can choose from them.

To formalize the approach, let $G=(V,E)$ be the county-level graph, $r \in V$ be a designated root county, and k' be a designated cluster size. For now, we seek a single connected cluster $S \subseteq V$ that contains r with population between Lk' and Uk' for which the size of the cut $\delta(S)=\{\{i,j\} \in E \mid |\{i,j\} \cap S|=1\}$ is minimum. The intent is that this cluster S will later be subdivided into k' districts, and its complement $V \setminus S$ will be subdivided into $k-k'$ districts. To promote this, we require $V \setminus S$ to be connected and to have a population between $L(k-k')$ and $U(k-k')$. Requiring $V \setminus S$ to be connected is not strictly required but is convenient for our purposes.

We introduce a binary assignment variable x_{ij} for each vertex $i \in V$ and each cluster number $j \in \{1,2\}$, which represent S and $V \setminus S$. We also introduce a binary variable y_e for each edge $e=\{u,v\} \in E$ indicating whether it is cut. The basic model, without contiguity constraints, is:

$$\min \sum_{e \in E} y_e \tag{1a}$$

$$\text{s.t. } x_{i1}+x_{i2}=1 \qquad \forall i \in V \tag{1b}$$

$$Lk' \leq \sum_{i \in V} p_i x_{i1} \leq Uk' \tag{1c}$$

$$L(k-k') \leq \sum_{i \in V} p_i x_{i2} \leq U(k-k') \tag{1d}$$

$$x_{i1}-x_{j1} \leq y_e \text{ and } x_{j1}-x_{i1} \leq y_e \qquad \forall e=\{i,j\} \in E \tag{1e}$$

$$x_{r1}=1 \tag{1f}$$

$$x,\ y \text{ binary.} \tag{1g}$$

The objective (1a) minimizes the number of cut edges. The assignment constraints (1b) ensure that each vertex is either assigned to S or its complement. Constraints (1c) and (1d) ensure population balance. Constraints (1e) ensure that if an edge is not cut, then its endpoints are either both assigned to S or neither is. Constraint (1f) forces the root r to be in S. Although this model could be simplified by replacing each instance of x_{i2} with $1-x_{i1}$, we prefer the presentation above for clarity and because MIP solvers will perform these substitutions in presolve anyway.

Now, consider the contiguity constraints. In our experience, the flow-based contiguity constraints of Shirabe (2005, 2009) work well for county-level instances, especially when the root is known *a priori*. In our case, we know that r will root S. So, we introduce a flow variable f_{ij} for each *directed* edge (i,j) and impose the following constraints, where $N(v)$ is the neighborhood of vertex v.

$$\sum_{j \in N(i)} (f_{ji}-f_{ij})=x_{i1} \qquad \forall i \in V \setminus \{r\} \tag{2a}$$

$$\sum_{j \in N(i)} f_{ji} \leq M x_{i1} \qquad \forall i \in V \setminus \{r\} \tag{2b}$$

$$f_{jr} = 0 \qquad \forall j \in N(r) \tag{2c}$$

$$f_{ij},\ f_{ji} \geq 0 \qquad \forall \{i,j\} \in E. \tag{2d}$$

Constraints (2a) ensure that each vertex selected in $S \setminus \{r\}$ consumes one unit of flow. The "big-M" constraints (2b) ensure that flow can only enter selected vertices, and we set $M = |V| - 1$. Constraints (2c) disallow flow from entering the root.

For the complement, we do not know a root *a priori*, and in this case separator inequalities (Carvajal et al. 2013, Wang et al. 2017, Oehrlein and Haunert 2017, Validi et al. 2022) work better as they are lightweight and do not introduce model symmetry. In our case, they take the form

$$x_{a2} + x_{b2} \leq 1 + \sum_{v \in R} x_{v2}, \tag{3}$$

where $a, b \in V$ are nonadjacent vertices that become disconnected when removing $R \subseteq V \setminus \{a, b\}$ from the graph. Because there are exponentially many of these inequalities, we implement them in a cut callback and use the algorithm of Fischetti et al. (2017) to find minimal violated inequalities.

The full model for finding a most compact cluster is then given by (1), (2), and (3). However, we seek not just one cluster, but the top t clusters. If using the Gurobi solver, one straightforward approach would be to change the `PoolSearchMode` parameter to enumerate the t best solutions. However, any given cluster can be paired with many different values of the f variables, possibly producing the same cluster over and over. Another approach would be to solve the model from scratch t times, each time adding a no-good cut for the previous values of the x variables. In principle, this would work, but duplicates effort. Instead, we record each solution ourselves in a callback and instruct the solver to keep searching by adding a no-good cut in the x space of variables. We terminate the search early (with the nonsensical cut $x_{r1} \leq -1$) when the best LP bound cannot beat any of the t current best solutions.

3.1. Case Study for Alabama

In the Supreme Court case *Allen v. Milligan* (2023), Alabama's congressional districts were challenged under Section 2 of the Voting Rights Act (VRA), which prohibits diluting the voting strength of protected minority groups. In particular, the enacted districts were criticized for dividing the state's Black Belt across multiple districts. An effect was that only one of the seven districts (14%) could elect Black voters' candidates of choice, even though more than 27% of the state is Black.

To bring a Section 2 lawsuit, the *Milligan* plaintiffs needed to satisfy the *Gingles* preconditions, which were established by the Supreme Court in *Thornburg v. Gingles* (1986). The first precondition to be shown is that the minority group is sufficiently numerous and geographically compact to constitute a majority in a single-member district, see also *Bartlett v. Strickland* (2009). That is, the plaintiffs must show that the minority group could achieve better representation in an alternative, reasonably configured map—in this case with *two* majority-Black districts, not one.

Evan Milligan himself was unable to draw such a map, and mathematics professor Moon Duchin was hired to draw demonstration districts (Buchanan 2023b). She used computer optimization techniques to draw preliminary maps, and then later drew four plans by hand. The fourth plan ("Plan D") had six county splits, which she testified to be minimum possible. Later, Alabama's attorney Edmund LaCour repeated the same claim to Justice Kavanaugh during Supreme Court oral arguments when complaining that Duchin's other plans were not reasonably configured, in part because they had more splits than the supposed minimum of six. The Supreme Court ruled in Milligan's favor, and a special master was later tasked with drawing remedial plans. The special master's plans "avoid[ed] county splits where possible", and all had at least six county splits.

Using our proposed integer programming techniques, we found the ten most compact county clusters of size two that are rooted at Jefferson County (whose seat Birmingham is nearly 70% Black). Among these ten clusters is the one from Figure 3. Meanwhile, its complement contains majority-Black cities such as Mobile, Montgomery, and Selma, as well as many counties from the Black Belt. By hand, we divided the cluster into two districts (with one being majority-Black) using one county split in Jefferson County. We then divided the cluster's complement into five districts (with one being majority-Black) using four county splits.

We arrive at the plan[2] in Figure 4, which has two majority-Black districts, 51.33% and 50.58% by voting age population (VAP). Importantly for *Gingles*, the districts are also reasonably configured; they are contiguous, satisfy a 1-person deviation, and have an average Polsby-Popper compactness score of 0.2211, which is comparable to that of Alabama's originally enacted plan (0.2203) and the Special Master's remedial plans (which were reported as 0.23, 0.24, 0.24). Also, contrary to impossibility claims made in expert testimony, cross examination, and Supreme Court oral arguments, the plan has just five county splits (and five precinct splits). We conclude that $k-1$ county splits should not be assumed minimum, even when constrained by contiguity, compactness, 1-person deviation, and minority representation.

3.2. Case Study for North Carolina

In the case *Harper v. Hall* (2022), the North Carolina Supreme Court overturned the state's enacted congressional districts for being an unconstitutional partisan gerrymander, with 10/14≈71.43% of the districts favoring Republicans despite the state's nearly even partisan makeup (49% vs. 48%). A

[2] `https://davesredistricting.org/join/7016a5b6-3bfe-46ec-b5e8-6010c26de508`

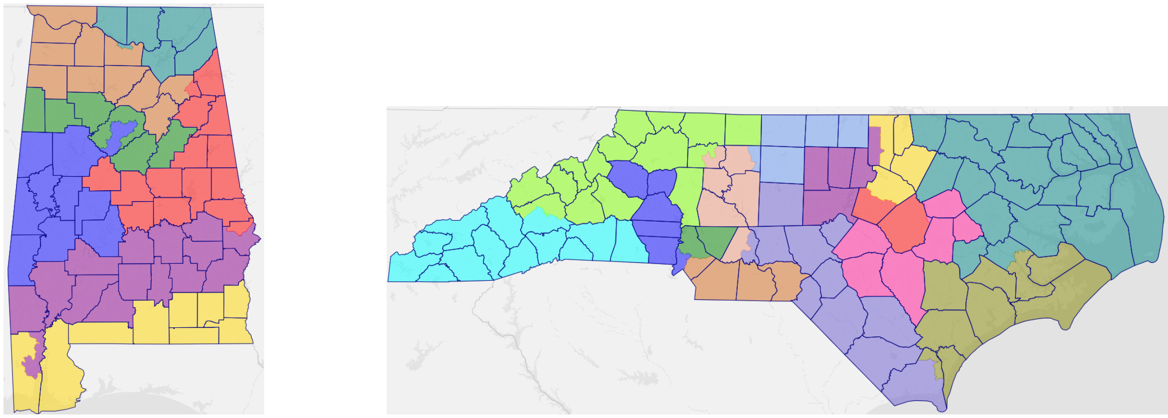

Figure 4 Reasonably configured plans for Alabama and North Carolina with fewer than $k-1$ county splits

remedial plan was drawn in which six districts favored Democrats, seven districts favored Republicans, and one was a tossup. The remedial plan had thirteen county splits (one less than the enacted plan). In 2023, after Republicans gained a majority on the North Carolina Supreme Court, the ruling was overturned, and the state's General Assembly enacted another Republican gerrymander that, at the time of writing, is the subject of several lawsuits.

The belief that $k-1$ is the minimum number of county splits also entered *Harper v. Hall*. The 2021 districting criteria adopted by North Carolina's House Committee on Redistricting and the Senate Committee on Redistricting and Elections state that "Division of counties in the 2021 Congressional plan shall only be made for reasons of equalizing population and consideration of double bunking" and that "VTDs [i.e., precincts] should be split only when necessary" (*Joint Meeting of Committees* 2021). Nevertheless, the enacted plan had 14 county splits and 25 precinct splits. These shortcomings were pointed out by expert witness and political science professor Jowei Chen (*Harper v. Hall* 2021), who wrote that "a congressional plan in North Carolina needs to contain only 13 county splits if the map-drawer is attempting to minimize the splitting of counties" and that "only 13 VTD splits" are necessary. He thus faulted the enacted plan for having "one more [county] split than is necessary" and "far more VTD splits than is necessary", violating the "mandated criteria" of "minimizing county splits [and] minimizing VTD splits". The belief that 13 splits is minimum continues to be repeated, e.g., by a conservative/libertarian foundation in North Carolina that criticized the way in which counties were split in the enacted plan (Jackson 2023).

Using our proposed integer programming techniques, we found the ten most compact county clusters of size three and four rooted at the two most populous counties: Mecklenburg County (which contains Charlotte) and Wake County (which contains Raleigh and Cary). Among these clusters, we pick one for Mecklenburg (C_1) and one for Wake (C_2) that are compatible, meaning that (C_1, C_2, C_3) is a county clustering, where $C_3 := C \setminus (C_1 \cup C_2)$, with the cluster sizes being $(k_1, k_2, k_3) = (3, 4, 7)$. This is the county clustering from Figure 3. By hand, we divided the clusters into three, four, and seven districts using two, three, and six county splits, respectively, giving a total of 11 county splits (and 11 precinct splits).

We arrive at the plan[3] in Figure 4. Among the 14 districts, five favor Democrats, six favor Republicans, and three are tossups. The plan fares well on various partisan fairness metrics, as reported by Dave's Redistricting App (2024), even though it was not optimized for them. For example, the plan has a mean-median score of −0.98%, which is substantially better (i.e., closer to zero) than the two gerrymanders (5.76% and 6.25%) and only slightly worse than the remedial plan (0.68%). Similar performance is observed for partisan bias (1.45%) compared to the gerrymanders (16.75% and 19.82%) and the remedial plan (0.25%). The districts are also reasonably configured; they are contiguous, satisfy a 1-person deviation, and have an average Polsby-Popper compactness score of 0.3329, which is better than the two gerrymanders (0.2974 and 0.2432) and the remedial plan (0.3234). Also, contrary to impossibility claims made in expert testimony, this plan has just 11 county splits (and 11 precinct splits). We conclude that $k-1$ county splits should not be assumed minimum, even when constrained by contiguity, compactness, 1-person deviation, and partisan fairness (or competitiveness).

4. Generating Whole-County Plans with Integer Programming

This section extends the approach to find contiguous, whole-county plans with 1-person deviation. We note that the direct application of an integer programming model is ill-suited for this task. Commercial MIP solvers will run for days on end without finding a feasible solution for instances like Iowa that have $k = 4$ districts when subjected to 1-person deviation (Shahmizad and Buchanan 2023). This poor performance persists regardless of which integer programming model is used: Hess (Hess et al. 1965) or labeling (Validi and Buchanan 2022); the manner in which contiguity is imposed: single-commodity flow (Hojny et al. 2021), multi-commodity flow (Shirabe 2009, Validi et al. 2022), separator constraints (Oehrlein and Haunert 2017, Validi et al. 2022); or the symmetry handling technique: diagonal-fixing (Validi and Buchanan 2022) or the extended formulation for partitioning orbitopes (Faenza and Kaibel 2009). We require a different approach.

At a high-level, the idea is to repeatedly carve a district from the state, much like the algorithm of McCartan and Imai (2023). One key difference is that McCartan and Imai aim to understand the *distribution* of possible plans, while we are interested in the tails, leading to differences in the carving strategy (randomized vs. optimization-minded).

Below, we propose a MIP-based districting heuristic. In it, $\mathcal{P}$ is a collection of partial plans (in which not all counties have been assigned to a district), and $\mathcal{C}$ is a collection of completed plans.

1. initialize $\mathcal{C} \leftarrow \{\}$ and $\mathcal{P} \leftarrow \{\{\}\}$

[3] `https://davesredistricting.org/join/061f823b-717f-4b61-aade-8d625d1b3001`

2. while $\mathcal{P} \neq \{\}$ do

- select and remove a partial plan P from $\mathcal{P}$
- let $V' = V \setminus (\cup_{D \in P} D)$ be the vertices unassigned in partial plan P
- let $k' = k - |P|$ be the number of unfinished districts in partial plan P
- if $k' = 1$, then add the completed plan $P \cup \{V'\}$ to $\mathcal{C}$ and continue
- pick a root county $r \in V'$
- using MIP techniques from Section 3, find (up to) t districts $\mathcal{D}$ in $G[V']$ that contain r
- for each district $D \in \mathcal{D}$, add new partial plan $P \cup \{D\}$ to $\mathcal{P}$

3. return $\mathcal{C}$

The heuristic initializes the collection of partial plans with a single empty plan (with all vertices being unassigned). Each iteration of the while loop extends a partial plan by one district. In it, a root vertex r is selected, and its top t most compact districts are found. For each of these districts, a new partial plan is obtained. If only $t' < t$ such districts exist (possibly $t' = 0$), then the heuristic only creates t' new partial plans from the given partial plan. Thus, the branching factor is at most t. In our implementation, the default value is $t = 10$, generating up to $10^{4-1} = 1000$ plans for Iowa. Generally, the heuristic can find up t^{k-1} plans, and the user can cast a wider net with larger t.

In principle, all plans could be found using a procedure like this, setting t to infinity. However, a better approach for enumerating *all* contiguous, whole-county plans (regardless of their population deviation or compactness properties) is the `enumpart` algorithm of Fifield et al. (2020), Kawahara et al. (2017), which is designed for this purpose. In particular, `enumpart` finds that Montana admits precisely 30,223 contiguous, whole-county plans with 1-person deviation. We thank Chris Kenny for carrying out this 3-day computation at our request (Kenny 2024). It should be noted, however, that the number of plans for Iowa is too huge to be enumerated in full; Fifield et al. (2020) resort to sampling 500 million plans, which they state is still "miniscule relative to the total number of valid partitions... into four districts, of which there are approximately 10^{24}".

4.1. Applying the MIP-Based Heuristic

This section applies the MIP-based districting heuristic to several US states, specifically Idaho, Iowa, Mississippi, Montana, Nebraska, and West Virginia. In each case, our approach finds *multiple* contiguous, whole-county congressional plans with 1-person deviation (or less).

We begin with Idaho, Montana, and West Virginia, which have two districts. Each admits contiguous, whole-county plans with 1-person deviation, see Figure 5. In fact, Idaho and West Virginia admit plans with 0-person deviation. West Virginia is particularly interesting, as it draws whole-county plans in practice. Its 2010 districts were upheld by the Supreme Court in *Tennant v. Jefferson County* (2012) in a *per curiam* opinion, despite exhibiting a 4871-person deviation, justified by the state's desires to keep counties whole and to preserve the cores of prior districts. Meanwhile, the current map has a 1582-person deviation. It is unclear to us whether it would survive a similar challenge, given that the deviation is still nontrivial and the cores of prior districts had to be disrupted after the state lost a seat in reapportionment.

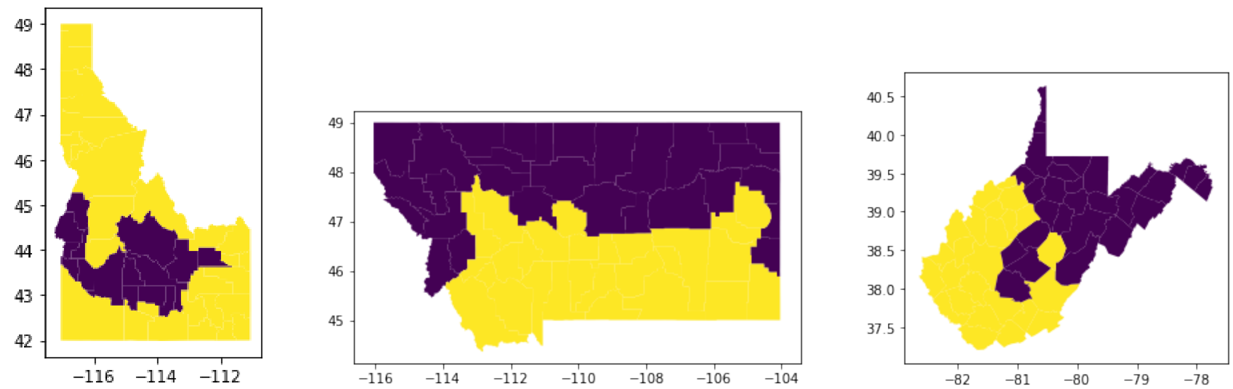

Figure 5 **Plans for Idaho, Montana, and West Virginia with zero county splits and 1-person deviation (or less)**

Next, we consider Nebraska, Mississippi, and Iowa, which have three or four districts. Each admits whole-county plans with 1-person deviation, see Figure 6. Iowa is a common test case for redistricting algorithms (Fifield et al. 2020, Becker and Solomon 2022, McCartan and Imai 2023, McCartan 2023) because it is the largest state that draws county-level plans in practice. Iowa was also the subject of a redistricting contest hosted by Dave Wasserman of the Cook Political Report in which the task was to find a contiguous, whole-county plan with the smallest deviation. The winner, Cory McCartan, used a carving strategy to find a plan with a 5-person deviation (Burger 2021). Meanwhile, our default implementation finds plans with *1-person deviation*; in fact, it finds 151 of them. By increasing the default parameter to $t = 25$, our implementation finds *1,104 such plans.*

We conclude that $k-1$ county splits should not be assumed minimum, even when constrained by contiguity and 1-person deviation. In fact, some states admit hundreds, thousands, or tens of thousands of such plans with *zero* county splits.

5. Conclusion

As we have seen, it is not unusual for a state to admit a districting plan with fewer than $k-1$ county splits, even when subjected to 1-person deviation. These counterexamples (or "accidental degeneracies" in the words of Nagle (2022)) are not rare flukes. This runs contrary to statements made by Autry et al. (2021) who wrote that "Given the extremely tight population constraints on congressional districts, it is reasonable to assume that there is no subset of counties that perfectly can accommodate a subset of the congressional districts." Not only do these county clusters exist,

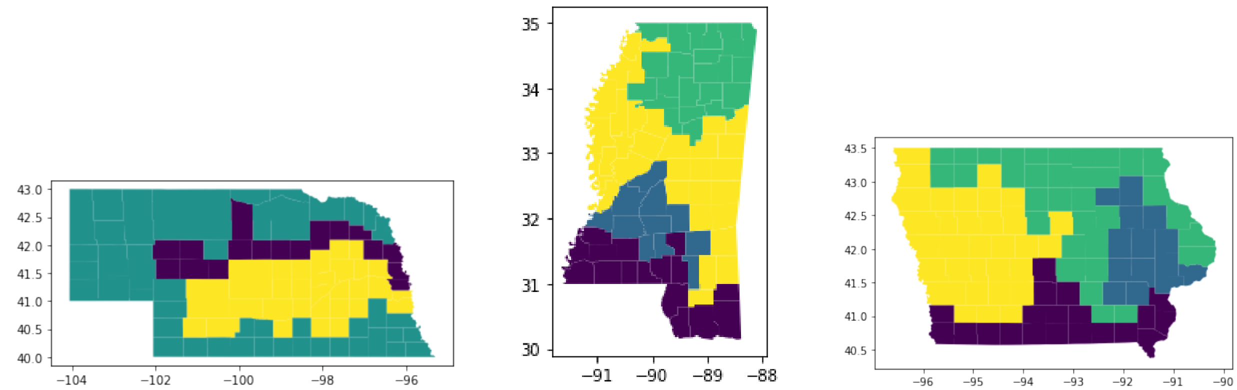

Figure 6 **Plans for Nebraska, Mississippi, and Iowa with zero county splits and 1-person deviation**

but in fact states like Iowa and Montana admit literally hundreds, thousands, or tens of thousands of *entire plans* satisfying 1-person deviation in which all counties are kept whole. This finding runs counter to many people's intuitions and can be chalked up to combinatorial explosion.

The optimization methods proposed in this paper can assist in the drawing of maps that simultaneously satisfy good government criteria (e.g., compactness, preservation of political subdivisions), minority representation, and partisan fairness. Indeed, our approach is flexible, providing mapmakers a "menu" of compact county clusters to choose from. Each cluster can be divided into districts however one chooses, either assisted by computer methods or by hand. To achieve fewer than $k-1$ county splits, the user need only to use two county clusters and to divide each cluster of size k' into districts using $k'-1$ county splits. The advances in MIP methodology proposed in this paper can assist mapmakers in this task.

To be clear, we make no normative claims about how many county splits is best in practice. It may well be that $k-1$ county splits (or more) can be justified when seeking to satisfy other criteria. Courts have also stated that they would like to avoid "county-split beauty contests" (*Allen v. Milligan* 2023). But, we should not treat a mathematical suspicion about splits as fact until it has been verified, nor should we confuse a normative belief with a fact about reality.

Future work may consider the tasks of finding the *minimum number of county splits* under 1-person deviation and finding the *most compact* plans under 1-person deviation. For states like North Carolina and Iowa, these problems seem to be extremely difficult. It is possible that set partitioning models, like those of Garfinkel and Nemhauser (1970) and Mehrotra et al. (1998), solved via branch-and-price, may be helpful in this regard. Our procedures for enumerating the top t most compact clusters may help in finding high-quality initial columns and in solving the pricing problem. However, preliminary experiments with this approach suggest that it would require a careful, nontrivial implementation, which is beyond the scope of the present paper.

References

Autry EA, Carter D, Herschlag GJ, Hunter Z, Mattingly JC (2021) Metropolized multiscale forest recombination for redistricting. *Multiscale Modeling & Simulation* 19(4):1885–1914.

Becker A, Duchin M, Gold D, Hirsch S (2021) Computational redistricting and the Voting Rights Act. *Election Law Journal: Rules, Politics, and Policy* 20(4):407–441.

Becker A, Gold D (2022) The gameability of redistricting criteria. *Journal of Computational Social Science* 5:1735–1777.

Becker A, Solomon J (2022) Redistricting algorithms. Duchin M, Walch O, eds., *Political Geometry: Rethinking Redistricting in the US with Math, Law, and Everything In Between* (Birkhauser).

Belotti P, Buchanan A, Ezazipour S (2023) Political districting to optimize the Polsby-Popper compactness score. Available on Optimization-Online.

Buchanan A (2021) PoolSearchMode with auxiliary variables (extended formulations)? `https://support.gurobi.com/hc/en-us/community/posts/360078219851-PoolSearchMode-with-auxiliary-variables-extended-formulations`.

Buchanan A (2023a) Political districting. Pardalos PM, Prokopyev OA, eds., *Encyclopedia of Optimization* (Cham: Springer International Publishing).

Buchanan A (2023b) Using optimization to support minority representation in Voting Rights Act Cases. *ORMS Today* 50(4):32–35.

Burger E (2021) The map whisperer: A conversation with Cory McCartan, winner of the 2021 Iowa redistricting challenge. *Little Village* URL `https://littlevillagemag.com/a-conversation-with-cory-mccartan-winner-of-the-2021-iowa-redistricting-challenge/`.

Carter D, Hunter Z, Teague D, Herschlag G, Mattingly J (2020) Optimal legislative county clustering in North Carolina. *Statistics and Public Policy* 7(1):19–29.

Carvajal R, Constantino M, Goycoolea M, Vielma JP, Weintraub A (2013) Imposing connectivity constraints in forest planning models. *Operations Research* 61(4):824–836.

Cervas J (2022) Report of the Special Master, *Harkenrider v. Hochul*. `http://jonathancervas.com/2022/NY/CERVAS-SM-NY-2022.pdf`.

Cervas JR, Grofman B (2020) Tools for identifying partisan gerrymandering with an application to congressional districting in Pennsylvania. *Political Geography* 76:102069.

Cohen-Addad V, Klein PN, Young NE (2018) Balanced centroidal power diagrams for redistricting. *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 389–396 (ACM).

Dave's Redistricting App (2024) Dave's Redistricting App. `https://davesredistricting.org/`, accessed: 2024-2-22.

DeFord D, Duchin M (2019) Redistricting reform in Virginia: Districting criteria in context. *Virginia Policy Review* 12(2):120–146.

DeFord D, Duchin M, Solomon J (2021) Recombination: A family of Markov chains for redistricting. *Harvard Data Science Review* 3(1).

DeFord D, Kimsey E, Zerr R (2023) Multi-balanced redistricting. *Journal of Computational Social Science* 6(2):923–941.

Dobbs KW, King DM, Jacobson SH (2023a) Redistricting optimization with recombination: A local search case study. *Computers & Operations Research* 160:106369.

Dobbs KW, Swamy R, King DM, Ludden IG, Jacobson SH (2023b) An optimization case study in analyzing Missouri redistricting. *INFORMS Journal on Applied Analytics* To appear.

Duchin M (2022) Explainer: Compactness, by the numbers. Duchin M, Walch O, eds., *Political Geometry*, 29–35 (Birkhauser).

Allen v Milligan (2021) Presentation of Alternative Congressional Districting Plans for Alabama. Available in the Supplemental Joint Appendix (3rd volume from April 25, 2022).

Allen v Milligan (2022a) Transcript of January 6, 2022 Preliminary Injunction Hearing. Available at `https://www.courtlistener.com/docket/61494291/105/2/milligan-v-allen/`.

Allen v Milligan (2022b) Transcript of Oral Arguments. Available at `https://www.supremecourt.gov/oral_arguments/argument_transcripts/2022/21-1086_1pd4.pdf`.

Allen v Milligan (2023) Report and Recommendation of the Special Master. Available at `https://thearp.org/documents/14352/AL_221-cv-1530_295.pdf`.

Harper v Hall (2021) Expert Report of Dr. Jowei Chen. Exhibit H, available at `https://static1.squarespace.com/static/5e909f4422f7a40a188de597/t/61a9588f318a824c40103bd3/1638488213466/Plaintiffs%27+PI+Motion+Exhibits+G-I.pdf`.

Joint Meeting of Committees (2021) Criteria Adopted by the Committees. Available at `https://webservices.ncleg.gov/ViewDocSiteFile/38467`.

Faenza Y, Kaibel V (2009) Extended formulations for packing and partitioning orbitopes. *Mathematics of Operations Research* 34(3):686–697.

Fifield B, Imai K, Kawahara J, Kenny CT (2020) The essential role of empirical validation in legislative redistricting simulation. *Statistics and Public Policy* 7(1):52–68.

Fischetti M, Leitner M, Ljubić I, Luipersbeck M, Monaci M, Resch M, Salvagnin D, Sinnl M (2017) Thinning out Steiner trees: a node-based model for uniform edge costs. *Mathematical Programming Computation* 9(2):203–229.

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

Gladkova T, Goldbloom-Helzner A, Khan M, Kolstoe B, Noory J, Schutzman Z, Stucky E, Weighill T (2019) Discussion of locality splitting measures. `https://github.com/vrdi/splitting/blob/master/SplittingReport.pdf`.

Guth L, Nieh A, Weighill T (2022) Three applications of entropy to gerrymandering. Duchin M, Walch O, eds., *Political Geometry*, 275–292 (Birkhauser).

Hebert JG, Vandenberg ME, Smith P (2010) *The Realist's Guide to Redistricting: Avoiding the Legal Pitfalls* (American Bar Association).

Hess S, Weaver J, Siegfeldt H, Whelan J, Zitlau P (1965) Nonpartisan political redistricting by computer. *Operations Research* 13(6):998–1006.

Hojny C, Joormann I, Lüthen H, Schmidt M (2021) Mixed-integer programming techniques for the connected max-k-cut problem. *Mathematical Programming Computation* 13(1):75–132.

Jackson A (2023) A redistricting tale of three counties split three ways. Available at `https://www.johnlocke.org/a-redistricting-tale-of-three-counties-split-three-ways/`.

Kaufman AR, King G, Komisarchik M (2021) How to measure legislative district compactness if you only know it when you see it. *American Journal of Political Science* 65(3):533–550.

Kawahara J, Inoue T, Iwashita H, Minato Si (2017) Frontier-based search for enumerating all constrained subgraphs with compressed representation. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences* 100(9):1773–1784.

Kenny C (2024) Montana counties with enumpart. `https://github.com/christopherkenny/mt-enumeration`.

Kenny CT, McCartan C, Simko T, Kuriwaki S, Imai K (2023) Widespread partisan gerrymandering mostly cancels nationally, but reduces electoral competition. *Proceedings of the National Academy of Sciences* To appear.

McCartan C (2023) Finding Pareto efficient redistricting plans with short bursts. *arXiv preprint arXiv:2304.00427* .

McCartan C, Imai K (2023) Sequential Monte Carlo for sampling balanced and compact redistricting plans. *Annals of Applied Statistics* 17(4):3300–3323.

McCartan C, Kenny CT, Simko T, Garcia III G, Wang K, Wu M, Kuriwaki S, Imai K (2022) Simulated redistricting plans for the analysis and evaluation of redistricting in the United States. *Scientific Data* 9(1):689.

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

Nagle J (2022) Euler's formula determines the minimum number of splits in maps of election districts. *Available at SSRN 4115039* .

NCSL (2020) 2010 redistricting deviation table. https://www.ncsl.org/redistricting-and-census/2010-redistricting-deviation-table, accessed: 2024-3-15.

NCSL (2021) Redistricting criteria. http://www.ncsl.org/research/redistricting/redistricting-criteria.aspx, accessed: 2024-2-22.

NCSL (2023) 2020 redistricting deviation table. https://www.ncsl.org/elections-and-campaigns/2020-redistricting-deviation-table, accessed: 2024-3-15.

Niemi RG, Grofman B, Carlucci C, Hofeller T (1990) Measuring compactness and the role of a compactness standard in a test for partisan and racial gerrymandering. *The Journal of Politics* 52(4):1155–1181.

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

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

Ramsay A (2022) Personal communication.

Shahmizad M, Buchanan A (2023) Political districting to minimize county splits. Available on Optimization-Online.

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.

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

Validi H, Buchanan A (2022) Political districting to minimize cut edges. *Mathematical Programming Computation* 14:623–672.

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

Wachspress J, Adler WT (2021) Split decisions: Guidance for measuring locality preservation in district maps. Center for Democracy & Technology, URL https://cdt.org/wp-content/uploads/2021/11/2021-11-04-Locality-splitting-report-final.pdf.

Wang HL (2022a) The 2020 census had big undercounts of Black people, Latinos and Native Americans. https://www.npr.org/2022/03/10/1083732104/2020-census-accuracy-undercount-overcount-data-quality, accessed: 2024-3-15.

Wang HL (2022b) These 14 states had significant miscounts in the 2020 census. https://www.npr.org/2022/05/19/1099810793/census-undercount-by-state-arkansas-florida-illinois-mississippi-tennessee-texas, accessed: 2024-3-15.

Wang Y, Buchanan A, Butenko S (2017) On imposing connectivity constraints in integer programs. *Mathematical Programming* 166(1-2):241–271.

Young HP (1988) Measuring the compactness of legislative districts. *Legislative Studies Quarterly* 13(1):105–115.

Zhang J, Validi H, Buchanan A, Hicks IV (2024) Linear-size formulations for connected planar graph partitioning and political districting. *Optimization Letters* 18(1):19–31.