

SEQUENTIAL MONTE CARLO FOR SAMPLING BALANCED AND COMPACT REDISTRICTING PLANS

BY CORY MCCARTAN^a AND KOSUKE IMAI^b

Department of Statistics, Harvard University, ^acmccartan@g.harvard.edu, ^bimai@harvard.edu

Random sampling of graph partitions under constraints has become a popular tool for evaluating legislative redistricting plans. Analysts detect partisan gerrymandering by comparing a proposed redistricting plan with an ensemble of sampled alternative plans. For successful application sampling methods must scale to maps with a moderate or large number of districts, incorporate realistic legal constraints, and accurately and efficiently sample from a selected target distribution. Unfortunately, most existing methods struggle in at least one of these areas. We present a new sequential Monte Carlo (SMC) algorithm that generates a sample of redistricting plans converging to a realistic target distribution. Because it draws many plans in parallel, the SMC algorithm can efficiently explore the relevant space of redistricting plans better than the existing Markov chain Monte Carlo (MCMC) algorithms that generate plans sequentially. Our algorithm can simultaneously incorporate several constraints commonly imposed in real-world redistricting problems, including equal population, compactness, and preservation of administrative boundaries. We validate the accuracy of the proposed algorithm by using a small map where all redistricting plans can be enumerated. We then apply the SMC algorithm to evaluate the partisan implications of several maps submitted by relevant parties in a recent high-profile redistricting case in the State of Pennsylvania. We find that the proposed algorithm converges faster and with fewer samples than a comparable MCMC algorithm. Open-source software is available for implementing the proposed methodology.

1. Introduction. In first-past-the-post electoral systems, legislative districts serve as the fundamental building block of democratic representation. In the United States, congressional redistricting, which redraws district boundaries in each state following the decennial census, plays a central role in influencing who is elected and hence what policies are eventually enacted. Because the stakes are so high, redistricting has been subject to intense political battles. Parties often engage in *gerrymandering* by manipulating district boundaries in order to amplify the voting power of some groups while diluting that of others.

In recent years the availability of granular data about individual voters has led to sophisticated partisan gerrymandering attempts that cannot be easily detected. At the same time, many scholars have focused their efforts on developing methods to uncover gerrymandering by comparing a proposed redistricting plan with a large collection of alternative plans that satisfy the relevant legal requirements. A primary advantage of such an approach, over the use of simple summary statistics, is its ability to account for the characteristics of each state's physical and political geography and state-specific redistricting rules.

For its successful application, a sampling algorithm for drawing alternative plans must: (1) be efficient enough to scale to maps with thousands of geographic units and a moderate or large number of districts, (2) simultaneously incorporate a variety of real-world legal constraints such as population balance (Section 3.1), geographical compactness (Section 3.3), and the preservation of administrative boundaries (Section 4.5), and (3) ensure these samples

are representative of a specific target population, against which a redistricting plan of interest can be evaluated. Although some have been used in several recent court challenges to existing redistricting plans, all existing algorithms run into limitations of varying severity with regards to at least one of these three key requirements.

Optimization-based (e.g., [Bozkaya, Erkut and Laporte \(2003\)](#), [Liu, Cho and Wang \(2016\)](#), [Macmillan \(2001\)](#), [Mehrotra, Johnson and Nemhauser \(1998\)](#)) and constructive Monte Carlo (e.g., [Chen and Rodden \(2013\)](#), [Cirincione, Darling and O'Rourke \(2000\)](#), [Magleby and Mosesson \(2018\)](#)) methods can be made scalable and incorporate many constraints. But they are not designed to sample from any specific target distribution. As a consequence, the resulting plans tend to differ systematically, for example, from a uniform distribution under certain constraints ([Cho and Liu \(2018\)](#), [Fifield et al. \(2020a, 2020b\)](#)). The absence of an explicit target distribution makes it difficult to interpret the ensembles generated by these methods and use them for statistical outlier analysis to detect gerrymandering.

MCMC algorithms (e.g., [Cannon et al. \(2022\)](#), [Carter et al. \(2019\)](#), [Chikina, Frieze and Pegden \(2017\)](#), [DeFord, Duchin and Solomon \(2021\)](#), [Fifield et al. \(2020a\)](#), [Mattingly and Vaughn \(2014\)](#), [Wu et al. \(2015\)](#)) can, in theory, sample from a specific target distribution and incorporate constraints through the use of an energy function. In practice, however, existing algorithms struggle to mix and traverse through a highly complex sampling space, making scalability difficult and accuracy hard to prove. Some of these algorithms make proposals by flipping precincts at the boundary of existing districts (e.g., [Fifield et al. \(2020a\)](#), [Mattingly and Vaughn \(2014\)](#)), rendering it difficult or even impossible to transition between points in the state space, especially as more constraints are imposed. More recent algorithms by [DeFord, Duchin and Solomon \(2021\)](#) and [Carter et al. \(2019\)](#) use spanning trees to make their proposals, and this has allowed these algorithms to yield greater moves and substantially improve mixing. Yet recent theoretical results suggest that even these larger moves may not be enough to traverse the entire state space and, therefore, may fail to converge to the correct distribution, if a realistic population balance constraint is imposed ([Akitaya et al. \(2022\)](#)).

We contribute to the ongoing scholarly efforts to address the above three key challenges by developing a new sequential Monte Carlo (SMC) algorithm, based on a similar but not identical spanning tree construction to [DeFord, Duchin and Solomon \(2021\)](#) and [Carter et al. \(2019\)](#) (see Sections 3 and 4). Like MCMC algorithms, the SMC algorithm generates samples which approximate the target distribution arbitrarily well as the sample size increases. But like constructive Monte Carlo methods, the SMC algorithm draws many separate plans from scratch rather than tweaking a single plan sequentially. This approach is better suited to the large discrete state space with a multimodal target distribution that characterizes redistricting problems. For example, in cases where existing MCMC proposals render the state space disconnected, the SMC algorithm can still converge to the target distribution. As we demonstrate in Sections 5 and 6, this sampling approach translates to faster convergence and smaller standard errors for a given computational budget. For larger and more complex redistricting sampling problems, the SMC algorithm can be easily parallelized to facilitate efficient computation.

The proposed algorithm proceeds by splitting off one district at a time, building up the redistricting plan piece by piece (see Figure 2 for an illustration). Each split is accomplished by drawing a spanning tree and removing one edge, which splits the spanning tree in two. We also extend the SMC algorithm so that it preserves administrative boundaries and certain geographical areas as much as possible, which is another common constraint considered in many real-world redistricting cases. An open-source software package, `redist`, is available for implementing the proposed algorithm ([Kenney et al. \(2020\)](#)).

The SMC algorithm is not without limitations. Like existing MCMC approaches, the SMC algorithm only guarantees convergence to the target distribution as the sample size

approaches infinity. Additionally, because the SMC algorithm involves repeated resampling with replacement, for a finite number of samples it can suffer from particle system collapse (Liu, Chen and Logvinenko (2001)), where many of the sampled plans share a small number of common districts (which originate as common ancestor particles in the SMC scheme). This can significantly increase sampling variability. Thus, it is important to understand the limitations of the proposed algorithm in finite samples, especially when dealing with large maps and many districts. In Section 4.4.4 we provide a set of diagnostics which can be used in practice to help identify if more samples are needed to reach convergence or if the constraints imposed by an analyst are too strong.

In Section 5 we validate the SMC algorithm using a small map for which all potential redistricting plans can be enumerated (Fifield et al. (2020b)). We demonstrate that the proposed algorithm samples accurately from a realistic target distribution and that our proposed diagnostics are a good proxy for total sampling error, which is generally unobservable. Section 6 applies the SMC algorithm to the 2011 Pennsylvania congressional redistricting and also compares its performance on this problem with an MCMC algorithm with the same transition kernel as the proposed SMC algorithm. We find that the SMC algorithm samples more efficiently than the MCMC approach. Section 7 concludes and discusses directions for future work.

2. The 2011 Pennsylvania congressional redistricting. We study the 2011 Pennsylvania congressional redistricting because it illustrates the salient features of the redistricting problem. We begin by briefly summarizing the background of this case and then explain the role of sampling algorithms used in the expert witness reports.

2.1. Background. Pennsylvania lost a seat in Congress during the reapportionment of the 435 U.S. House seats following the 2010 Census. In Pennsylvania the General Assembly, which is the state's legislative body, draws new congressional districts, subject to gubernatorial veto. At the time the General Assembly was controlled by Republicans, and Tom Corbett, also a Republican, served as governor. In the 2012 election, which took place under the newly adopted 2011 districting map, Democrats won 5 seats while Republicans took the remaining 13. Under the previous plan, the split was 7–12.

In June 2017, the League of Women Voters of Pennsylvania filed a lawsuit alleging that the 2011 plan adopted by the Republican legislature violated the state constitution by diluting the political power of Democratic voters. The case worked its way through the state court system, and on January 22, 2018, the Pennsylvania Supreme Court issued its ruling, writing that the 2011 plan “clearly, plainly and palpably violates the Constitution of the Commonwealth of Pennsylvania, and, on that sole basis, we hereby strike it as unconstitutional.” (League of Women Voters v. Commonwealth (2018)).

The court ordered that the General Assembly adopt a remedial plan and submit it to the governor, who would in turn submit it to the court by February 15, 2018. In its ruling the court laid out specific requirements that had to be satisfied by all proposed plans:

composed of compact and contiguous territory; as nearly equal in population as practicable; and which do not divide any county, city, incorporated town, borough, township, or ward, except where necessary to ensure equality of population.

The leaders of the Republican Party in the General Assembly drew a new map, but the Democratic governor, Tom Wolf, refused to submit it to the court, claiming that it, too, was an unconstitutional gerrymander. Instead, the court received remedial plans from seven parties: the petitioners, the League of Women Voters, the respondents, the Republican leaders of the General Assembly, the governor—a Democrat; the lieutenant governor—also a Democrat, the Democratic Pennsylvania House minority leadership, the Democratic Pennsylvania Senate

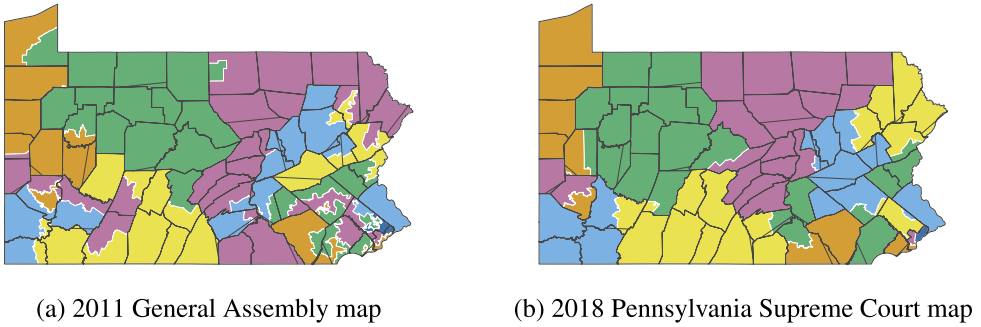


FIG. 1. Comparison of the 2011 map drawn by the General Assembly and the final map imposed by the Supreme court in 2018. County lines are shown in dark gray, and district boundaries that do not coincide with county boundaries are in white.

minority leadership, and the intervenors, which included Republican party candidates and officials. Ultimately, the Supreme Court drew its own plan and adopted it on February 19, 2018, arguing that it was “superior or comparable to all plans submitted by the parties.” Figure 1 shows the remedial plan created by the Supreme Court as well as the 2011 map adopted by the General Assembly, which were found on the court’s case page.

The constraints explicitly laid out by the court as well as the numerous remedial plans submitted by the parties make the 2011 Pennsylvania redistricting a useful case study that evaluates redistricting plans.

2.2. The role of sampling algorithms. The original finding that the 2011 General Assembly plan was a partisan gerrymander was, in part, based on different outlier analyses performed by two academic researchers, Jowei Chen and Wesley Pegden, who served as the petitioner’s expert witnesses. Chen randomly generated two sets of 500 redistricting plans according to a constructive Monte Carlo algorithm based on [Chen and Rodden \(2013\)](#). He considered population balance, contiguity, compactness, avoiding county and municipal splits, and, in the second set of 500, avoiding plans that placed more than two incumbents in the same district (at least one pair of incumbents in the same district was necessary, given that Pennsylvania lost a seat from 2000 to 2010). Pegden ran a reversible Markov chain similar to that used in the MCMC algorithm of [Mattingly and Vaughn \(2014\)](#) for one trillion steps and computed upper bounds of p -values using the method of [Chikina, Frieze and Pegden \(2017\)](#). This method was also used in a follow-up analysis by Moon Duchin, who served as an expert for Governor Wolf ([Duchin \(2018\)](#)). Both petitioner experts concluded that the 2011 plan was an extreme outlier according to compactness, county and municipal splits, and the number of Republican and Democratic seats implied by statewide election results.

The respondents also retained an expert academic witness, Wendy Tam Cho, who directly addressed the sampling-based analyses of Chen and Pegden. Cho criticized Chen’s analysis for not sampling from a specified target distribution. She also criticized Pegden’s analysis by arguing that his Markov chain only made local explorations of the space of redistricting plans and could not, therefore, have generated a representative sample of all valid plans, though the p -values computed using the [Chikina, Frieze and Pegden \(2017\)](#) method explicitly do not require mixing of the Markov chain (see also [Cho and Rubinstein-Salzedo \(2019\)](#), and [Chikina, Frieze and Pegden \(2019\)](#)). We do not directly examine the intellectual merits of the specific arguments put forth by the expert witnesses. However, these methodological debates are also relevant for other cases where simulation algorithms have been extensively used by expert witnesses (e.g., [Rucho v. Common Cause \(2019\)](#); [Common Cause v. Lewis \(2019\)](#); [Covington v. North Carolina \(2017\)](#); [Harper v. Lewis \(2020\)](#)) and highlight the difficulties in practically applying existing sampling algorithms to actual redistricting problems.

First, the distributions that some of these algorithms sample from are not made explicit, leaving open the possibility that the generated ensemble is systematically different from the true set of all valid plans. Second, even when the distribution is known, MCMC algorithms used to sample from it may be prohibitively slow to mix and cannot yield a representative sample. These challenges motivate us to design an algorithm that accurately samples from a specific target distribution and incorporates most common redistricting constraints, while being efficient and scalable.

3. Sampling balanced and compact districts.

3.1. *The setup.* Redistricting plans are ultimately aggregations of geographic units such as counties, voting precincts, or census blocks. The usual requirement that the districts in a plan be contiguous necessitates consideration of the spatial relationship between these units. The natural mathematical structure for this consideration is a graph $G = (V, E)$, where $V = \{v_1, v_2, \dots, v_m\}$ consists of m nodes, representing the geographic units of redistricting, and E contains edges connecting units which are legally adjacent.

A labeled redistricting plan on G consists of n districts, where each district is a collection of nodes. A labeled plan is described by a function $\xi : V \rightarrow \{1, 2, \dots, n\}$, where $\xi(v) = i$ implies that node v is in district i . In practice, we will be interested in *unlabeled* plans, since the assignment of sets of precincts to labels is arbitrary and has no impact on the real-world aspects of the district such as its population, demographic composition, or partisan lean. Define an equivalence relation by $\xi_1 \cong \xi_2$ if there exists a permutation σ such that $\xi_1(v) = \sigma(\xi_2(v))$ for all v . Then an unlabeled redistricting plan can be viewed as an equivalence class under \cong , denoted $[\xi]$; nothing in what follows will depend on the particular labeled plan representative ξ .

We let $V_i(\xi)$ and $E_i(\xi)$ denote the nodes and edges contained in district i under a given redistricting plan ξ , so $G_i(\xi) = (V_i(\xi), E_i(\xi))$ represents the induced subgraph that corresponds to district i under the plan. We suppress the dependence on ξ , when it is clear from context, writing $G_i = (V_i, E_i)$. Since each node belongs to only one district, we have $V = \bigcup_{i=1}^n V_i(\xi)$ and $V_i(\xi) \cap V_{i'}(\xi) = \emptyset$ for any redistricting plan ξ . In addition, we require that each district be contiguous, that is, that G_i is a connected graph for all i .

Beyond connectedness redistricting plans are always required to have roughly equal population in every district. To formalize this requirement, let $\text{pop}(v)$ denote the population of node v . Then the population of a district i may be written as $\text{pop}(V_i(\xi)) := \sum_{v \in V_i(\xi)} \text{pop}(v)$. We quantify the discrepancy between a given plan and the ideal of equal population in every district by the *maximum population deviation*,

$$\text{dev}(\xi) := \max_{1 \leq i \leq n} \left| \frac{\text{pop}(V_i)}{\text{pop}(V)/n} - 1 \right|,$$

where $\text{pop}(V)$ is the total population. Some courts and states have imposed hard maximums on this quantity, for example, $\text{dev}(\xi) \leq D = 0.05$ for state legislative redistricting ([National Conference of State Legislatures \(2021\)](#)).

The proposed algorithm samples plans by way of spanning trees on each district, that is, subgraphs of $G_i(\xi)$, which contain all vertices, no cycles, and are connected. Let T_i represent a spanning tree for district i whose vertices and edges are given by $V_i(\xi)$ and a subset of $E_i(\xi)$, respectively. The collection of spanning trees from all districts together form a spanning forest. Each node belongs to one spanning tree in the forest, and this assignment corresponds to a redistricting plan. However, a single redistricting plan may correspond to multiple spanning forests because each district may admit more than one spanning tree.

For a given redistricting plan, we can compute the exact number of spanning forests in polynomial time using the determinant of a submatrix of the graph Laplacian, according to

the Matrix Tree Theorem of Kirchhoff (see [Tutte \(1984\)](#)). Thus, for a graph H , if we let $\tau(H)$ denote the number of spanning trees on the graph, we can represent the number of spanning forests that correspond to a redistricting plan ξ as $\tau(\xi) := \prod_{i=1}^n \tau(G_i(\xi))$. This fact will play an important role in the definition of our sampling algorithm and its target distribution.

3.2. The target distribution. The algorithm is designed to sample an unlabeled plan $[\xi]$ with probability

$$(1) \quad \pi([\xi]) \propto \exp\{-J(\xi)\} \tau(\xi)^\rho \mathbf{1}_{\{\xi \text{ connected}\}} \mathbf{1}_{\{\text{dev}(\xi) \leq D\}},$$

where the indicator functions ensure that the plans meet population balance and connectedness criteria, $\tau(\xi)$ measures the compactness of the districts in ξ (see Section 3.3), and J encodes additional constraints on the types of plans preferred. As done in Section 6, we often use a reasonably strict population constraint such as $D = 0.001$ and $D = 0.005$. The parameter $\rho \in \mathbb{R}_0^+$ is chosen to control the compactness of the generated plans.

This target distribution π has both substantive and theoretical justifications. First, it incorporates two constraints which are almost always present in real-world redistricting: the support of π is restricted to contiguous plans and those that meet a population deviation threshold. Second, it represents the unique maximum entropy distribution on the set of redistricting plans satisfying these two universal constraints and the moment conditions implied by the other constraints, that is, $\mathbb{E}_\pi[\log \tau(\xi)] = \mu_\tau$ and $\mathbb{E}_\pi[J(\xi)] = \mu_J$ for some constants μ_τ and μ_J (see [Cover and Thomas \(\(2006\), Theorem. 12.1.1, originally of Boltzmann\)](#)).

Thus, our target distribution ensures that all plans meet contiguity and population requirements and, *on average*, satisfy a compactness standard as well as any other additional constraints (through the function J). It is not surprise, therefore, that this class of target distributions has been used by other work developing redistricting sampling algorithms ([Fifield et al. \(2020a\)](#), [Herschlag, Ravier and Mattingly \(2017\)](#)).

The generality of the additional constraint function J is intentional, as its exact form and number imposed on the redistricting process varies by state and by the type of districts being drawn; any type of constraint may be incorporated by choosing a J which is small for preferred plans and large otherwise. For example, a preference for plans close to an existing plan ξ_{sq} may be encoded as

$$(2) \quad \begin{aligned} J_{sq}(\xi) &= -\frac{\beta}{\log n} \text{VI}(\xi, \xi_{sq}) \\ &:= -\frac{\beta}{2 \log n} \sum_{i,j=1}^n \frac{P_{ij}}{\text{pop}(V)} \left(\log\left(\frac{P_{ij}}{\text{pop}(V_j(\xi_{sq}))}\right) + \log\left(\frac{P_{ij}}{\text{pop}(V_j(\xi))}\right) \right), \end{aligned}$$

where $\beta \in \mathbb{R}^+$ controls the strength of the constraint and $P_{ij} = \text{pop}(V_i(\xi) \cap V_j(\xi_{sq}))$ is the population shared between district i of ξ and j of ξ_{sq} . The function $\text{VI}(\cdot, \cdot)$ represents the variation of information (also known as the shared information distance), which is the difference between the joint entropy and the mutual information of the distribution of population over the new districts ξ relative to the existing districts ξ_{sq} ([Cover and Thomas \(2006\)](#)). When ξ is any relabeling of ξ_{sq} , then $J_{sq}(\xi) = 0$. In contrast, when ξ evenly splits the nodes of each district of ξ_{sq} among the districts of ξ , then $J_{sq}(\xi) = \beta$. This distance measure will prove useful later in quantifying the diversity of a sample of redistricting plans (see also [Guth, Nieh and Weighill \(2022\)](#)).

There exist other formulations of constraints and considerations in choosing a set of weights that balance constraints against each other (see, e.g., [Bangia et al. \(2017\)](#), [Fifield et al. \(2020a\)](#), [Herschlag, Ravier and Mattingly \(2017\)](#)). Here we focus on sampling from the broad class of distributions characterized by equation (1), which have been used in other

work; we do not address the important but separate problem of picking a specific instance of this class for a given redistricting problem.

The flexibility of J can be deceptive, however. The algorithm operates efficiently only when the additional constraints imposed by J are not too severe. Even a small number of strong constraints incorporated into J can dramatically limit the number of valid plans and considerably complicate the process of sampling (Chatterjee and Diaconis (2018)). The Markov chain algorithms developed to date partially avoid this problem by moving toward maps with lower J over a number of steps but, in general, including more constraints makes it even more difficult to transition between valid redistricting plans. Approaches, such as simulated annealing (Bangia et al. (2017), Herschlag, Ravier and Mattingly (2017)) and parallel tempering (Fifield et al. (2020a)), have been proposed to handle multiple constraints, but these can be difficult to calibrate in practice and provide few, if any, theoretical guarantees.

In practice, we usually find that the most stringent constraints are those involving population deviation, compactness, and administrative boundary splits. As shown later, we address this issue by designing our algorithm to directly satisfy these constraints. Weak additional constraints do not generally have a substantial effect on the sampling efficiency, though there are exceptions. Monitoring the distribution of the weights and the overall sampling efficiency is crucial to obtaining a good sample, as we discuss later.

3.3. Spanning forests and compactness. One common redistricting requirement is that districts be geographically compact, though nearly every state leaves this term undefined. Dozens of numerical compactness measures have been proposed, with the Polsby–Popper score (Polsby and Popper (1991)) perhaps the most popular. Defined as the ratio of a district’s area to that of a circle with the same perimeter as the district, the Polsby–Popper score is constrained to $[0, 1]$, with higher scores indicating more compactness.

Other scholars have proposed a graph-theoretic measure known as *edge-cut compactness* (Dube and Clark (2016), DeFord, Duchin and Solomon (2021)). This measure counts the number of edges that must be removed from the original graph to partition it according to a given plan. Formally, it is defined as

$$\text{rem}(\xi) := 1 - \frac{\sum_{i=1}^n |E_i(\xi)|}{|E(G)|},$$

where we have normalized to the total number of edges.

Plans that involve cutting many edges will necessarily have long internal boundaries, driving up their average district perimeter (and driving down their Polsby–Popper scores), while plans that cut as few edges as possible will have relatively short internal boundaries and much more compact districts. Additionally, given the high density of voting units in urban areas, plans which cut fewer edges will tend to avoid drawing district lines through the heart of these urban areas. This has the welcome side effect of avoiding splitting cities and towns, and in doing so helping to preserve “communities of interest,” another common redistricting consideration.

Empirically, this graph-based compactness measure is highly correlated with $\log \tau(G) - \log \tau(\xi)$ (we often observe a correlation in excess of 0.99). It is difficult to precisely characterize this relationship, except in special cases, because $\tau(\xi)$ is calculated as a matrix determinant (McKay (1981)). However, this quantity is strongly controlled by the product of the degrees of each node in the graph, $\prod_{i=1}^m \deg(v_i)$ (Kostochka (1995)). Removing an edge from a graph decreases the degree of the vertices at either end by one, so we would expect $\log \tau(G)$ to change by approximately $2\{\log \bar{d} - \log(\bar{d} - 1)\}$ with this edge removal, where \bar{d} is the average degree of the graph. This implies a linear relation $\log \tau(G) - \log \tau(\xi) \approx \text{rem}(\xi) \cdot 2\{\log \bar{d} - \log(\bar{d} - 1)\}$, and hence $\tau(\xi)^\rho \approx C_1 \exp(-C_2 \rho \text{rem}(\xi))$, where C_1 and C_2 are some constants depending on the details of the map.

As a result, a greater value of ρ in the target distribution corresponds to a preference for fewer edge cuts and, therefore, a redistricting plan with more compact districts. This and the considerations given in the literature (DeFord, Duchin and Solomon (2021), Dube and Clark (2016)) suggest that the target distribution in equation (1) with $\rho = 1$ (or another positive value) is a good choice for sampling compact districts. The choice of $\rho = 1$ is computationally convenient, as it allows us to avoid calculating $\tau(\xi)$ as part of sampling (an asymptotic bottleneck) and yet, usually, produces satisfactorily compact districts. Of course, if another compactness metric is desired, one can simply set $\rho = 0$ and incorporate the alternative metric into J . This will preserve the algorithm’s efficiency to the extent that the alternative metric correlates with the edge-removal measure of compactness. Setting $\rho = 0$ by itself, however, will make sampling intractable in most cases, just as choosing an extreme J will.

4. The proposed algorithm. The proposed algorithm samples redistricting plans by sequentially drawing districts over $n - 1$ iterations of a splitting procedure. This is fundamentally different from existing MCMC approaches, which change an existing plan according to some transition kernel. The iterations of the proposed algorithm are from district to district within a single plan, whereas the iterations in an MCMC algorithm are from plan to plan.

Our algorithm begins by partitioning the original graph $G = (V, E) = (\tilde{V}_0, \tilde{E}_0) = \tilde{G}_0$ into two induced subgraphs: $G_1 = (V_1, E_1)$, which will constitute a district in the final map, and the remainder of the graph $\tilde{G}_1 = (\tilde{V}_1, \tilde{E}_1)$, where $\tilde{V}_1 = V \setminus V_1$ and \tilde{E}_1 consists of all the edges between vertices in \tilde{V}_1 . Next, the algorithm takes \tilde{G}_1 as an input graph and partitions it into two induced subgraphs, one which will become a district G_2 and the remaining graph \tilde{G}_2 . The algorithm repeats the same splitting procedure until the final $(n - 1)$ th iteration whose two resulting partitions, G_{n-1} and $\tilde{G}_{n-1} = G_n$, become the final two districts of the redistricting plan.

Figure 2 provides an illustration of this sequential procedure. To sample a large number of redistricting plans from the target distribution, given in equation (1), at each iteration, the algorithm samples many candidate partitions, discards those which fail to meet the population constraint, and then resamples a certain number of the remainder according to importance weights, using the resampled partitions at the next iteration. The rest of this section explains the details of the proposed algorithm.

4.1. The splitting procedure. We first describe the splitting procedure, which is similar to the merge-split Markov chain proposals of DeFord, Duchin and Solomon (2021) and Carter et al. (2019). It proceeds by drawing a random spanning tree T , identifying the k_i most promising edges to cut within the tree, and selecting one such edge at random to create two induced subgraphs. Spanning trees are an attractive way to split districts, as the removal of a single edge induces a partition with two connected components, and spanning trees can be sampled uniformly (Wilson (1996)).




FIG. 2. The sequential splitting procedure applied to the state of Iowa, where four congressional districts are created at the county level.

Algorithm 1 Splitting procedure to generate one district

Input: Initial graph \tilde{G}_{i-1} and a parameter $k_i \in \mathbb{Z}^+$:

(a) Draw a single spanning tree T on \tilde{G}_{i-1} uniformly from the set of all such trees using Wilson’s algorithm.

(b) Each edge $e \in E(T)$ divides T into two components, $T_e^{(1)}$ and $T_e^{(2)}$. For each edge, compute the following population deviation for the two districts that would be induced by cutting T at e :

$$d_e^{(1)} = \left| \frac{\sum_{v \in T_e^{(1)}} \text{pop}(v)}{\text{pop}(V)/n} - 1 \right| \quad \text{and} \quad d_e^{(2)} = \left| \frac{\sum_{v \in T_e^{(2)}} \text{pop}(v)}{\text{pop}(V)/n} - 1 \right|.$$

Let $d_e = \min\{d_e^{(1)}, d_e^{(2)}\}$, and index the edges in ascending order by this quantity so that we have $d_{e_1} \leq d_{e_2} \leq \dots \leq d_{e_{m_i}}$, where $m_i = |\tilde{V}_{i-1}|$.

(c) Select one edge e^* uniformly from the top k_i edges, $\{e_1, e_2, \dots, e_{k_i}\}$, and remove it from T , creating a spanning forest $(T_{e^*}^{(1)}, T_{e^*}^{(2)})$ which induces a partition $(G_i^{(1)}, G_i^{(2)})$.

(d) If $d_{e^*}^{(1)} \leq d_{e^*}^{(2)}$, that is, if $T_{e^*}^{(1)}$ induces a district that is closer to the optimal population than $T_{e^*}^{(2)}$ does, set $G_i = G_i^{(1)}$ and $\tilde{G}_i = G_i^{(2)}$; otherwise, set $G_i = G_i^{(2)}$ and $\tilde{G}_i = G_i^{(1)}$.

As part of the full sampling procedure (Algorithm 2), after splitting, we check that the population of the new district G_i falls within the bounds $[P_i^-, P_i^+]$, where

$$P_i^- = \max \frac{\text{pop}(V)}{n} (1 - D), \quad \text{pop}(\tilde{V}_{i-1}) - \frac{n-i}{n} \text{pop}(V) (1 + D) \quad \text{and}$$

$$P_i^+ = \min \frac{\text{pop}(V)}{n} (1 + D), \quad \text{pop}(\tilde{V}_{i-1}) - \frac{n-i}{n} \text{pop}(V) (1 - D).$$

These bounds also ensure that it will be possible for future iterations to generate valid districts out of \tilde{G}_i . If $\text{pop}(V_i) \notin [P_i^-, P_i^+]$, then the entire redistricting plan is rejected, and the sampling process begins again. While the rate of rejection varies by map and by iteration, we generally encounter acceptance rates at each iteration between 5% and 30%, which are not so low as to make sampling from large maps intractable. Algorithm 1 details the steps of the splitting procedure, where at the first iteration we take $\tilde{G}_0 = G$.

4.2. The sampling probability. The above sequential splitting procedure does not generate plans from the target distribution π . We denote the sampling measure by q and write the sampling probability for a given connected plan ξ at iteration i as $q(G_i | \tilde{G}_{i-1})$, since each new district G_i depends only on the leftover map area \tilde{G}_{i-1} from the previous iteration. This probability can be written as the probability that we cut an edge along the boundary of the new district, integrated over all spanning trees which could be cut to form the district, that is,

$$(3) \quad q(G_i | \tilde{G}_{i-1}) = \sum_{T \in \mathcal{T}(\tilde{G}_{i-1})} q(G_i | T) \tau(\tilde{G}_{i-1})^{-1},$$

where $\mathcal{T}(\cdot)$ represents the set of all spanning trees of a given graph and we have relied on the fact that Wilson’s algorithm draws spanning trees uniformly.

The key is that, for certain choices of k_i (the number of edges considered to be cut at iteration i), the probability that an edge is cut is independent of the trees that are drawn. Let $\text{ok}(T)$ represent the number of edges on any spanning tree T that induce balanced partitions with population deviation below D , that is,

$$\text{ok}(T) := |\{e \in E(T) : d_e \leq D\}|.$$

Then define $K_i := \max_{T \in \mathcal{T}(\tilde{G}_{i-1})} \text{ok}(T)$, the maximum number of such edges across all spanning trees. Furthermore, let $\mathcal{C}(G, H)$ represent the set of edges joining nodes in a subgraph G to nodes in a subgraph H . We have the following result for the splitting probability for new districts whose populations lie inside the bounds defined above (see Appendix A in the Supplementary Material (McCartan and Imai (2023)) for the proof).

LEMMA 4.1. *The probability of splitting a valid new district G_i from an existing area \tilde{G}_{i-1} , using Algorithm 1 with parameter $k_i \geq K_i$, is*

$$(4) \quad q(G_i | \tilde{G}_{i-1}, \text{pop}(V_i) \in [P_i^-, P_i^+]) = \frac{\tau(G_i)\tau(\tilde{G}_i)}{\tau(\tilde{G}_{i-1})k_i} |\mathcal{C}(G_i, \tilde{G}_i)|.$$

4.3. *Sequential Monte Carlo.* We follow a sequential Monte Carlo approach (Doucet, de Freitas and Gordon (2001), Liu, Chen and Logvinenko (2001)) to generate draws from the target distribution rather than simply performing $n - 1$ iterations of Algorithm 1 and resampling or reweighting at the final stage. A sequential approach is also useful in operationalizing the rejection procedure to enforce the population constraint.

The proposed procedure is presented as Algorithm 2. The algorithm is governed by a parameter $\alpha \in (0, 1]$, which has no effect on the target distribution nor the asymptotic accuracy of the algorithm. Rather, α may be adjusted to maximize the efficiency of sampling. To generate S redistricting plans, at each iteration of the splitting procedure $i \in \{1, 2, \dots, n - 1\}$, we resample and split the existing plans one at a time, rejecting those which do not meet the

Algorithm 2 Sequential Monte Carlo (SMC) Algorithm

Input: Graph G to be split into n districts, target distribution parameters $\rho \in \mathbb{R}_0^+$ and constraint function J , and sampling parameters $\alpha \in (0, 1]$ and $k_i \in \mathbb{Z}^+$, with $i \in \{1, 2, \dots, n - 1\}$:

(a) Generate an initial set of S plans $\{\tilde{G}_0^{(1)}, \tilde{G}_0^{(2)}, \dots, \tilde{G}_0^{(S)}\}$ and corresponding weights $\{w_0^{(1)}, w_0^{(2)}, \dots, w_0^{(S)}\}$, where each $\tilde{G}_0^{(j)} := G$ and $w_0^{(j)} = 1$.

(b) For each splitting iteration $i \in \{1, 2, \dots, n - 1\}$:

(1) Until there are S valid plans:

(i) Sample a partial plan \tilde{G}_{i-1} from $\{\tilde{G}_{i-1}^{(1)}, \tilde{G}_{i-1}^{(2)}, \dots, \tilde{G}_{i-1}^{(S)}\}$ according to weights $(\prod_{l=1}^{i-1} w_l^{(j)})^\alpha$.

(ii) Split off a new district from \tilde{G}_{i-1} through one iteration of the splitting procedure (Algorithm 1), creating a new plan (G_i, \tilde{G}_i) .

(iii) If the newly sampled plan (G_i, \tilde{G}_i) satisfies $\text{pop}(V_i) \in [P_i^-, P_i^+]$, save it; otherwise, reject it.

(2) Calculate weights for each of the new plans $w_i^{(j)} = \frac{\tau(G_i^{(j)})^{\rho-1}}{|\mathcal{C}(G_i^{(j)}, \tilde{G}_i^{(j)})|}$.

(c) Calculate final weights

$$(5) \quad w^{(j)} = \exp\{-J(\xi^{(j)})\} \psi([\xi^{(j)}])^{-1} \left(\prod_{i=1}^{n-2} w_i^{(j)}\right)^{(1-\alpha)} w_{n-1}^{(j)} (\tau(G_n^{(j)}))^{\rho-1}.$$

(d) Output the S final plans $\{\xi^{(j)}\}_{j=1}^S$, where $\xi^{(j)} = (G_1^{(j)}, \dots, G_{n-1}^{(j)}, G_n^{(j)})$, and the final weights $\{w^{(j)}\}_{j=1}^S$.

population constraints, until we obtain S new plans for the next iteration. This rejection process can be viewed as a form of partial rejection control (Liu, Chen and Logvinenko (2001), Liu, Chen and Wong (1998)) or a version of the AliveSMC algorithm (LeGland and Oudjane (2005), Peters, Fan and Sisson (2012)).

The weights at each stage serve three purposes. The first is to account for the extraneous $|\mathcal{C}(G_i, \tilde{G}_i)|$ term, which appears in the splitting probability $q(G_i | \tilde{G}_{i-1}, \text{pop}(V_i) \in [P_i^-, P_i^+])$ but not in the target distribution π . The second is to account for differences in the compactness parameter ρ ; the splitting procedure generates plans according to $\rho = 1$, but this may be different from the target value. The third purpose is to adjust for the imbalances between the labeled plans, generated by the splitting procedure, and the unlabeled plans, which are the target of sampling. While there are $n!$ labeled plans corresponding to every unlabeled plan, not every labeled plan may be sampled according to the sequential procedure. In fact, the number of labeled plans which are in the support of the SMC proposal distribution varies from one unlabeled plan to another. It is this imbalance which necessitates an additional correction.

For a labeled plan ξ , let G/ξ denote the district-level quotient graph corresponding to the plan; that is, the nodes are districts $\{1, \dots, n\}$, and edges connect adjacent districts. Notice that the splitting procedure ensures that the leftover area after each split, \tilde{G}_i , is a connected graph. This means that, for every $1 \leq i \leq n - 1$, the subgraph of G/ξ induced by the vertices $\{i + 1, i + 2, \dots, n\}$ (the districts drawn after split i) is connected. We call any labeled plan ξ a *sequentially valid labeling* if this property holds and denote by $\psi([\xi])$ the number of sequentially valid labelings corresponding to an unlabeled plan $[\xi]$. The SMC weights incorporate $\psi([\xi])$ as a correction term. Calculation of $\psi([\xi])$ is discussed below in Section 4.4.2.

From the output of Algorithm 2, one last resampling of S plans using the final weights can be performed to generate a final sample. Alternatively, the weights can be used directly to estimate the expectation of some statistics of interest under the target distribution, that is, $H = \mathbb{E}_\pi(h(\xi))$, using the self-normalized importance sampling estimate $\hat{H} = \sum_{j=1}^S h(\xi^{(j)})w^{(j)} / \sum_{j=1}^S w^{(j)}$.

The sampled plans are not completely independent, because the weights in each step must be normalized before resampling and because the resampling itself introduces some dependence. Precisely quantifying the amount of dependence is difficult. However, as we demonstrate in Section 5, the dependence is not large enough to cause measurable bias in summary statistics of interest.

It is difficult to precisely characterize the computational complexity of the entire SMC algorithm since the rejection sampling introduces a random component, which depends on the difficulty of sampling a new district within the population bounds. This random complexity is also shared by existing MCMC approaches, which must redraw proposals if they are invalid. Appendix C analyzes the complexity of generating each sample without accounting for the rejection step. For both SMC and MCMC samplers, the total sampling time increases linearly in the sample size. In contrast with MCMC approaches, however, step (b)(1) of Algorithm 2 can be embarrassingly parallelized, since each resample-split-check requires interaction only with the previous set of samples.

The weights in the proposed algorithm are chosen to match existing general SMC algorithms with partial rejection control. These existing algorithms provide guarantees as to the convergence of the samples to the target distribution. One such result, which will suffice for our purposes, is the following central limit theorem.

PROPOSITION 4.2. *Let $\pi_S = \sum_{j=1}^S w^{(j)} \delta_{[\xi^{(j)}]}$ be the weighted particle approximation generated by Algorithm 2. Then for all measurable h on unlabeled plans, as $S \rightarrow \infty$,*

$$\sqrt{S}(\mathbb{E}_{\pi_S}[h([\xi])] - \mathbb{E}_\pi[h([\xi])]) \xrightarrow{d} \mathcal{N}(0, V_{\text{SMC}}(h))$$

for some asymptotic variance $V_{\text{SMC}}(h)$.

A proof is given in Appendix A. This central limit theorem implies consistency (in S) of any quantity derived from the weighted samples. However, since this convergence is in probability (w.r.t. the algorithm’s sampling probability), the proposition does not establish that $\pi_S \xrightarrow{d} \pi$ almost surely. While the almost sure convergence result exists for a standard SMC algorithm (Del Moral, Doucet and Jasra (2006)), we do not know of an extension to the case of partial rejection control.

In some cases the constraints incorporated into $J(\xi)$ admit a natural decomposition to the district level as $\prod_{i=1}^n J'(G_i)$, for example, a preference for districts which split as few counties as possible or against districts which would pair off incumbents. In these cases an extra term of $\exp\{-J'(G_i^{(j)})\}$ can be added to the weights $w_i^{(j)}$ in each stage, and the same term can be dropped from the final weights $w^{(j)}$. This can be particularly useful for more stringent constraints; incorporating J' in each stage allows the importance resampling to “steer” the set of redistricting plans toward those which are preferred by the constraints. This same idea can be used to amortize the contribution of $\psi([\xi])$ to the final weights over the preceding SMC iterations; details can be found in Appendix C.

4.4. *Practical implementation and use.* The SMC algorithm described above allows for an accurate characterization of the target distribution given in equation (1) as the sample size goes to infinity. In practice, of course, we must use a finite number of samples. Additionally, the algorithm relies on choosing k_i , which can be challenging since K_i is typically unknown, and on calculating $\psi([\xi])$, which has no closed-form expression.

This section discusses these practical implementation challenges as well as selection of the parameter α . We also describe the use of diagnostics (Section 4.4.4) that can help identify when the algorithm is or is not performing well and, therefore, if more samples or different constraints are required. Further details about implementation may be found in Appendix C.

4.4.1. *Choosing k_i .* The accuracy of the algorithm is theoretically guaranteed only when the number of edges considered for removal at each stage is at least the maximum number of edges across all graphs which induce districts G_i with $\text{dev}(G_i) \leq D$, that is, $k_i \geq \tilde{K}_i$. Unfortunately, K_i is generally unknown in practice. We could conservatively set $k_i = |\tilde{V}_{i-1}| - 1$, the number of edges in each spanning tree, but such a choice results in a prohibitively inefficient algorithm—the random edge selected for removal will with high probability induce an invalid partition, leading to a rejection of the entire map. In practice, we estimate k_i before each SMC stage by generating random spanning trees, as we detail in Appendix C. Theoretical support for this estimate is given by Proposition C.1, which bounds the inaccuracy of the approximation with a user-selectable parameter.

4.4.2. *Calculating $\psi([\xi])$.* The number of sequentially valid labelings $\psi([\xi])$ corresponding to an unlabeled plan $[\xi]$ can vary significantly across unlabeled plans and has no closed-form expression. Unfortunately, $\psi([\xi])$ grows rapidly in n but not nearly as fast as $n!$, the total number of labelings per unlabeled plan, which makes both direct and approximate calculation methods challenging. We adopt a hybrid approach in practice. When $n \leq 13$, we directly compute $\psi([\xi])$ with a recursive divide-and-conquer algorithm. When $n > 13$, we can estimate $\psi([\xi])$ to arbitrary precision using a particular importance sampling scheme. Both of these approaches, which are detailed in Appendix C, do not increase the runtime of the SMC algorithm appreciably.

4.4.3. *Choosing α .* As noted above, as long as $\alpha \in (0, 1]$ the SMC algorithm is asymptotically valid. Larger values are more aggressive in downweighting unlikely plans (those which are overrepresented in q vs. π), which may lead to less diversity in the final sample,

while smaller values of α are less aggressive, which can result in more variable final weights and more wasted samples. Liu, Chen and Logvinenko (2001) recommend a default choice of $\alpha = 0.5$, which we have found appropriate, though, in general, we have not found the algorithm's performance to be very sensitive to α within the range $0.5 \leq \alpha \leq 0.9$.

4.4.4. *Diagnostics.* As with any complex sampler, application of the SMC algorithm in practice can be greatly facilitated by a set of diagnostic measures. Diagnostics can never prove that an algorithm is working correctly, but they can identify situations when the algorithm fails and shed light on why. We discuss several useful diagnostics below, all of which are implemented in our open-source software.

Our primary recommendation is to perform multiple independent runs of the SMC algorithm, which enables both checking for nonconvergence to the target distribution as well as the calculation of standard errors for quantities of interest. For convergence, we adopt the Gelman–Rubin \hat{R} statistic (Gelman and Rubin (1992)), which compares between-run variation to within-run variation. If the latter is an appreciable fraction of the former, then independent runs produce different results, and the algorithm has not converged. For our purposes we will consider an algorithm not to have converged if \hat{R} values for statistics of interest exceed 1.05, though higher or lower thresholds may be appropriate for given computational budgets and other applied considerations. If the algorithm has not converged for a particular S , analysts should run the algorithm again with a larger S , until \hat{R} and other diagnostics, discussed below, indicate that the results are trustworthy. At this point the samples from the multiple runs can be combined to increase the precision of the estimates (Gelman et al. ((2014), Chapter 11)).

We use the rank-normalized and folded \hat{R} of Vehtari et al. (2021), which makes the statistic more robust to heavy tails and more sensitive to discrepancies in scale, not just location. For our purposes \hat{R} is advantageous, since it is computed on summary statistics and thus provides a measure of convergence tailored to practical quantities of interest. Additionally, it is unitless and so can be compared across runs and indeed across algorithms, as we will do in Section 5.

Multiple runs can also be easily used to compute standard errors of expectations taken relative to the sampled plans. Recently, Lee and Whiteley (2018) and Olsson and Douc (2019), among others, have developed a method to estimate SMC variance by using information on the ancestry of each sampled particle (here, plan). While we have implemented this method in our open-source software and found it to be correct on average, the resulting estimates tend to be noisier than those obtained from a simple multiple-run standard error.

Other diagnostics are useful for assessing the overall quality of the sample and pinpointing where issues arise when the algorithm is not working well. One measure of quality is *sample diversity*, or how different the samples are from each other, which we measure using the variation of information metric shown in equation (2). Similar plans will have a variation of information near zero, while plans which are extremely different will have a variation of information closer to 1. A nondiverse sample will have many copies of similar or identical plans, which tends to increase sampling error and reduces the interpretability of the generated samples.

When sampling problems arise, a closer look at each iteration of Algorithm 2 can help identify the cause. A low acceptance rate in step (b)(1)(3), or weights which are highly variable or have a heavy right tail in step (b)(2), can lower the sampling efficiency and diversity. Examining the variance of the log-weights at each iteration can be useful. Variances which increase consistently across iterations can be a sign to use a higher α , which will sample more aggressively earlier on to avoid a build-up of variable weights. We have also found that the number of unique plans from the previous iteration which survive to the subsequent iteration is a highly useful indicator of sampling problems, since it captures inefficiencies introduced

by both the rejection and resampling steps. If the number of unique plans is substantially below what would be expected from a uniform sample with replacement from the same population, then some kind of bottleneck or other sampling problem is generally present.

All of the diagnostics presented here are sample statistics, in that they vary across runs of the random SMC algorithm. If the number of independent parallel runs were increased to infinity, these diagnostics would converge to a “population value” for the particular sampling problem and choice of algorithm parameters. But with a finite number of runs, there will be variation around this population value, which can lead to diagnostic values that are too optimistic or pessimistic. No diagnostic is foolproof. However, taken together, the set of diagnostics presented here is designed to catch insufficient sample sizes and too-strong constraints with high probability. Indeed, these diagnostics formalize and extend the so-called “multistart heuristic” advocated by some MCMC redistricting practitioners (Cannon et al. (2022)).

4.5. *Incorporating administrative boundary constraints.* Another common requirement for redistricting plans is that districts “to the greatest extent possible” follow existing administrative boundaries such as county and municipality lines.¹ In theory, this constraint can be formulated using a J function which penalizes maps for every county line crossed by a district. In practice, however, we can more efficiently generate desired maps by directly incorporating this constraint into our sampling algorithm.

Fortunately, with a small modification to the proposed algorithm, we can sample redistricting plans proportional to a similar target distribution but with the additional constraint that the number of administrative splits not exceed $n - 1$. The hard constraint of $n - 1$ splits cannot be adjusted upward or downward, unfortunately, but a preference for even fewer administrative splits may be incorporated through the J function.

Let A be the set of administrative units, such as counties. We can relate these units to the nodes by way of a labeling function $\eta : V \rightarrow A$ that assigns each node to its corresponding unit. This function induces an equivalence relation \sim_η on nodes, where $v \sim_\eta u$ for nodes v and u iff $\eta(v) = \eta(u)$. If we quotient G by this relation, we obtain the administrative-level multigraph G / \sim_η , where each vertex is an administrative unit and every edge corresponds to an edge in G which connects two nodes in different administrative units. We can write the number of administrative splits as

$$\text{spl}(\xi) = \left(\sum_{a \in A} \sum_{i=1}^n C(\eta^{-1}(a) \cap \xi^{-1}(i)) \right) - |A|,$$

where $C(\cdot)$ counts the number of connected components in the subgraph $\eta^{-1}(a) \cap \xi^{-1}(i)$.

To implement this constraint, we draw the spanning trees in step (a) of Algorithm 2 in two substeps such that we sample from a specific subset of all spanning trees. First, we use Wilson’s algorithm to draw a spanning tree on each administrative unit $a \in A$, and then we connect these spanning trees to each other by drawing a spanning tree on the quotient multigraph \tilde{G}_i / \sim_η . Figure 3 illustrates this process. This approach is similar to the independently-developed multiscale merge-split algorithm of Autry et al. (2020).

Drawing the spanning trees in two steps limits the trees used to those which, when restricted to the nodes $\eta^{-1}(a)$ in any administrative unit a , are still spanning trees. The importance of this restriction is that cutting any edge in such a tree will either split the map exactly along administrative boundaries (if the edge is on the quotient multigraph) or split

¹If a redistricting plan must always respect these boundaries, we can simply treat the administrative units as the nodes of the original graph to be partitioned.

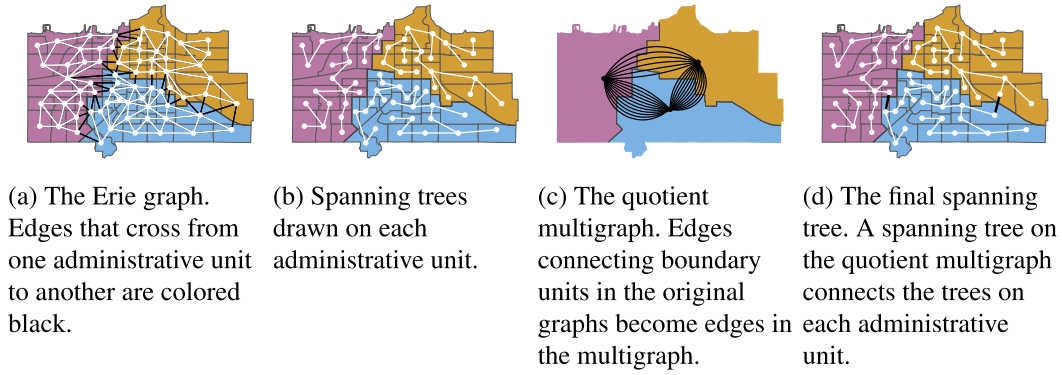


FIG. 3. The two-step spanning tree sampling procedure applied to the city of Erie, Pennsylvania, with three arbitrary administrative units indicated by the colored sections of each map.

one administrative unit in two and preserve administrative boundaries everywhere else. Since the algorithm has $n - 1$ stages, this limits the support of the sampling distribution to maps with no more than $n - 1$ administrative splits.

The two-step construction makes clear that the total number of such spanning trees is given by

$$(6) \quad \tau_\eta(\tilde{G}_i) = \tau(\tilde{G}_i / \sim_\eta) \prod_{a \in A} \tau(\tilde{G}_i \cap \eta^{-1}(a)),$$

where $\tilde{G}_i \cap \eta^{-1}(a)$ denotes the subgraph of \tilde{G}_i , which lies in unit a , and we take $\tau(\emptyset) = 1$. Replacing τ with τ_η in the expression for the weights $w_i^{(j)}$ and $w^{(j)}$ then gives the modified algorithm that asymptotically samples from

$$(7) \quad \pi_\eta([\xi]) \propto \exp\{-J(\xi)\} \tau_\eta(\xi)^\rho \mathbf{1}_{\{\xi \text{ connected}\}} \mathbf{1}_{\{\text{dev}(\xi) \leq D\}} \mathbf{1}_{\{\text{spl}(\xi) \leq n-1\}}.$$

This idea can, in fact, be extended to arbitrary levels of nested administrative hierarchy. We can, for example, limit not only the number of split counties but also the number of split cities and census tracts to $n - 1$ each, since tracts are nested within cities, which are nested within counties. To do so, we begin by drawing spanning trees using Wilson’s algorithm on the smallest administrative units. We then connect spanning trees into larger and larger trees by drawing spanning trees on the quotient graphs of each higher administrative level. (In this approach cities which cross county boundaries must be split in two, and regions outside of cities must be treated as their own pseudo-city.) Even with multiple levels of administrative hierarchy, the calculation of the number of spanning trees is still straightforward, by analogy to equation (6).

5. An empirical validation study. Although the proposed algorithm has desirable theoretical properties, it is important to empirically assess its performance (Fifield et al. (2020b)). We examine whether or not the proposed algorithm can produce a sample of redistricting maps that is actually representative of a target distribution and how the diagnostics of Section 4.4.4 correlate with sampling accuracy. Replication code and data for this analysis and all other analyses are available in the Supplementary Material (McCartan and Imai (2023)).

Our validation setting is the six-by-six grid map shown in Figure 4. This map is small enough to obtain all possible redistricting plans with six contiguous districts using an efficient enumeration algorithm (see, e.g., Fifield et al. (2020b)). Each precinct in the grid has an equal population. There are over 356 billion unlabeled plans with six districts, 451,206 of




FIG. 4. The six-by-six map used in validation (left). The remaining panels compare the distribution of the number of removed edges in the enumerated distribution (purple bars) and the sampled distribution (grey) over a number of sample sizes S . The dots are the mean histogram estimate across the 24 independent runs and represent the algorithm’s bias. The vertical lines are 90% confidence intervals for the histogram estimates and represent the sampling variability. They are calculated from the standard deviation of the estimates across the runs.

which have districts with exactly balanced populations.² We demonstrate that the proposed algorithm can efficiently approximate a realistic target distribution on this set of balanced plans. Appendix B contains an additional validation study with four districts using a 50-precinct map taken from the state of Florida, where the population of each precinct varies.

We set the target distribution by choosing $\rho = 1$, which, as discussed above, will generate compact districts. Since this value of ρ makes the proposal and sampling distributions as close as possible, the SMC algorithm will operate at peak efficiency. The validation study in Appendix B explores a wider range of values of ρ . As we have noted in Section 3, values of ρ other than 1 may not be computationally feasible for larger redistricting problems with many more precincts.

We run Algorithm 2 on a logarithmically-spaced range of sample sizes S ranging from 10 to 10,000. While $S = 10$ is far too small to use in any real analysis, we include it here for illustrative purposes to demonstrate the high variability and potential bias of the algorithm when S is too small. At each sample size, we perform 24 independent runs of the algorithm, each of size S , which allows for accurate calculation of standard errors and \hat{R} for summary statistics.

To validate the algorithm, we compare the distribution of a summary statistics from each sample to the true distribution obtained by reweighting the enumerated plans according to the target measure, which is proportional to $\tau(\xi)$. The summary statistic used here is the number of edges that must be removed to form each set of districts, $\text{rem}(\xi) \cdot |E(G)|$. As discussed in Section 3.3, this statistic is highly correlated with the target $\tau(\xi)$ and so is a good test of the algorithm.

These comparisons are summarized in the right three panels of Figure 4, which show the true distribution of the number of edges removed in purple, with the estimated histogram overlaid as solid points, for a range of sample sizes. Each histogram estimate produced by the SMC algorithm is accompanied by a 90% confidence interval produced using the standard errors calculated across the 24 independent runs (i.e., the standard deviations of the estimates for each run). With as few as $S = 100$ samples, the SMC estimates are already close to the true distribution. In fact, if we tested the null hypothesis that each estimate is centered at the true value, we would only reject for one of the 12 nonzero true values at the 10% level, within the range of what would be expected by chance given the sampling variability. As the sample size increases, the variance of the SMC estimates decreases; until by $S = 10,000$ samples, it is mostly negligible. Overall, the agreement between the sampling and target distributions

²The enumerated plans are publicly available at <https://github.com/gerrymandr/trees/tree/main/enumerations>.

is excellent, indicating that the SMC algorithm is able to accurately sample from the target distribution in this setting.

We next study more quantitatively the quality of the SMC sample as the sample size increases and, in particular, how well the diagnostics recommended in Section 4.4.4 can be used as a proxy for this quality. To do so, we focus on three particular estimands that can be obtained from the distribution of edges removed: the probability of removing exactly 22 edges (the target median, true probability 0.2328), the probability of removing exactly 18 edges (true value 0.0542), and the probability of removing more than 26 edges (true value 0.0204). The latter two statistics are especially useful as tests of the algorithm’s sampling accuracy, since they require estimating small tail probabilities which may correspond to very few re-districting plans. On the lower tail, there are just two possible unlabeled plans which remove exactly 18 edges, though together they comprise over 5% of the target distribution. The upper tail probability corresponds to the p -value a researcher would use for testing whether an enacted plan with 26 removed edges was significantly less compact than would be expected under the target distribution.

For each sample size from the SMC algorithm, we calculate the \hat{R} diagnostic and standard errors for these three summary statistics. Note that neither of these calculations requires access to the enumerated ground truth. We also compute the bias and the root mean-square error (RMSE) for the estimates, comparing to the true values from the enumerated target distribution. Ideally, the always-computable \hat{R} and standard errors will well represent the generally uncomputable bias and RMSE.

Figure 5 summarizes the results of this experiment. As expected, as the SMC sample size S increases, \hat{R} , the standard errors, and the relative RMSE also decrease. The \hat{R} values fall below our recommended cutoff of 1.05 by $S = 60$, indicating that for samples this large and greater, each independent run is producing similar estimates and the SMC algorithm has likely converged to the target distribution. Indeed, the estimates’ bias is closely centered around zero in this range. However, standard errors are significantly larger for the smaller sample sizes. This tracks the pattern observed in the validation above: at small sample sizes, the SMC estimates were, on average, unbiased for the target distribution but had significant sampling variability.

Importantly, the calculated standard errors track the actual RMSE extremely closely, demonstrating their value as an observable proxy for this measure of sampling accuracy.




FIG. 5. Convergence diagnostic \hat{R} , standard errors, bias, and RMSE for the same three compactness summary statistics, calculated across 24 independent runs of the SMC algorithms. Values are plotted versus sample size S per run.

In other words, with low \hat{R} indicating likely algorithmic convergence (and, therefore, unbiasedness), the overall algorithmic error is driven almost entirely by the sampling variability, which is measurable with the standard errors.

6. Analysis of the 2011 Pennsylvania redistricting. As discussed in Section 2, in the process of determining a remedial redistricting plan to replace the 2011 General Assembly map, the Pennsylvania Supreme Court received submissions from seven parties. In this section we compare four of these maps to both the original 2011 plan and the remedial plan ultimately adopted by the court. We study the governor’s plan and the House Democrats’ plan, the petitioner’s plan (specifically, their “Map A”), which was selected from an ensemble of 500 plans used as part of the litigation, and the respondent’s plan, which was drawn by Republican officials.

6.1. *The setup.* To evaluate these six plans, we drew reference maps from the target distribution given in equation (7) by using the proposed algorithm along with the modifications presented in Section 4.5 to cap the number of split counties at 17 (out of a total of 67), in line with the court’s mandate. We set $\rho = 1$, to put most of the sample’s mass on compact districts, and enforced $\text{dev}(\xi) \leq 0.001$ to reflect the “one person, one vote” requirement.

This population constraint translates to a tolerance of around 700 people, in a state where the median precinct has 1121 people. Like most research on redistricting, we use precincts because they represent the smallest geographical units for which election results are available. To draw from a stricter population constraint we would need to use the 421,545 census blocks in Pennsylvania rather than the 9256 precincts, which would significantly increase the computational burden.

6.2. *Comparison with an analogous MCMC algorithm.* We first compare the computational efficiency of the SMC algorithm with a merge-split-type MCMC algorithm (Carter et al. (2019), DeFord, Duchin and Solomon (2021)). The merge-split algorithm here uses a spanning tree-based proposal identical to the splitting procedure described in Algorithm 1: it merges adjacent districts, draws a spanning tree on the merged district, and splits it to ensure the population constraint is met. A Metropolis–Hastings correction is used, analogous to the SMC reweighting step. In fact, the implementation (contained in our open-source `redist` package) shares much of the same code base, which allows for performance comparisons that better reflect the different algorithmic strategies rather than different algorithmic implementations.³

The MCMC chains are each initialized with a random SMC sample and run for 500 warm-up iterations before samples are recorded. We let the sample size vary from 100 to 4000 for the SMC algorithm and from 100 to 8000 for the MCMC algorithm. For each sample size, we perform four independent runs of each algorithm. The average acceptance rate for the MCMC algorithm across the four chains was 88.7%.

Figure 6(a) shows the \hat{R} values of two summary statistics (Democratic seats and compactness) by sample size S per run. By this convergence heuristic, the SMC algorithm converges more quickly and with lower \hat{R} values than the MCMC algorithm for a given sample size. According to the check that $\hat{R} \leq 1.05$, the SMC algorithm has approximately converged with $S = 1600$, while the MCMC algorithm has still not converged after with than 8000 post-warmup samples (over 7000 unique accepted proposals), at which point the MCMC algorithm has been running 2.4 times longer than the SMC algorithm.

³A particularly performant implementation of a similar algorithm (Reversible ReCom, Cannon et al. (2022)) may be found at <https://github.com/pjrulc/frcw.rs>.

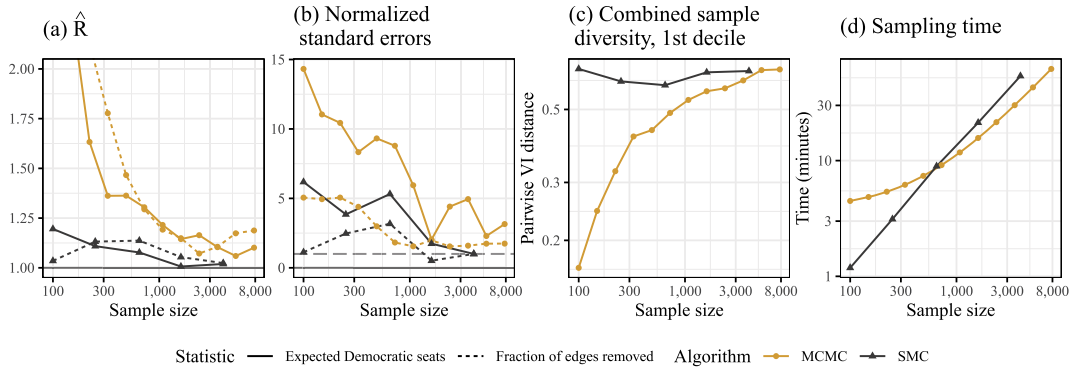


FIG. 6. A comparison of sampling efficiency for SMC and MCMC algorithms using the same transition kernel: (a) Convergence diagnostic \hat{R} for two summary statistics plotted against the sample size. (b) Standard errors for two statistics, normalized by the standard error for the SMC algorithm at 4000 samples (to put the two statistics on the same scale), by sample size. (c) First decile of the pairwise variation of information distance (a measure of sample diversity) by sample size. Distance measures were computed with the combined sample of four independent runs. (d) Sampling time by sample size.

We find a similar story in examining the standard errors of the two statistics, which are shown in Figure 6(b) after being normalized to be on comparable scales. With four independent runs per algorithm, standard error estimates themselves are rather noisy, but the overall trend mirrors the findings for \hat{R} . Even at large sample sizes, the MCMC standard errors are 1.5 to five times larger than their SMC counterparts, although the MCMC standard errors are not meaningful without the underlying chains having converged.

Figure 6(c) examines the sampling efficiency from the perspective of sample diversity. As discussed in Section 4.4.4, the pairwise variation of information (VI) distance provides a way to quantify how similar the plans in each sample are to one another. If the distribution of these pairwise VI distances has a long lower tail, then there are a significant number of plans which are very similar. Figure 6(c) shows the first decile of the pairwise VI distances in the combined sample (i.e., pooled across all four runs) for each algorithm as the sample size increases. Because of the sequential nature of the MCMC algorithms, the sample diversity doesn't match that of the SMC algorithms until there are at least 3000 samples. This is reflected in the lag-1 autocorrelation of the MCMC sample, which was 0.982 for Democratic seats and 0.984 for compactness. Autocorrelation is, of course, not a direct measure of sampling efficiency, but this comparison serves to underscore the differences in the types of samples generated by each algorithm for finite sample sizes.

Figure 6(d) shows how long each algorithm takes to run for a given number of samples. Because of the fixed warm-up cost of the MCMC algorithm, it is slower than the SMC algorithm for fewer than 600 samples. Asymptotically, both algorithms have runtime linear in the number of samples. But the per-sample cost of the MCMC algorithm is lower, since only two districts are redrawn rather than the entire map.

6.3. Compactness and county splits. While the results of the comparison above suggest that 1600 SMC samples is enough for Pennsylvania, we perform the rest of our analysis with 4000 samples from one of the SMC runs. Apart from the \hat{R} diagnostic, examined in detail above, the other SMC diagnostics of Section 4.4.4 all indicated no problems with the generated sample.

Figures 7(a) and (b) show distribution of the fraction of edges removed $\text{rem}(\xi)$ and the number of county splits $\text{spl}(\xi)$ across the reference maps generated by our algorithm. The figure also shows these values for each of the six plans. The 2011 General Assembly plan is a




FIG. 7. Summary statistics for the sampled and comparison plans: (a) compactness (smaller means more compact) and (b) county splits. Panel (c) shows the Democratic two-party vote share by district, where within each plan districts are ordered by Democratic vote share.

clear outlier for both statistics, being far less compact and splitting far more counties than any of the reference plans and all of the remedial plans. Among the remedial plans, the petitioner’s is almost an outlier in being *too* compact, although this is perhaps not surprising—the map was generated by an algorithm explicitly designed to optimize over criteria such as population balance and compactness (Chen (2017)). The other plans’ compactness falls within the range of the reference sample.

Figure 7(b) shows that all of the submitted remedial plans as well as the reference sample followed the court’s instructions to limit the number of county splits. Yet around half of the reference maps split fewer than 13 counties, the minimum number of splits found in any of the remedial plans. This may be a result of the strict population constraint (all six plans were within one person of equal population across all districts) or a different prioritization between the various constraints imposed.

6.4. *Partisan analysis.* While important, the outlier status of the General Assembly plan, as regards compactness and county splits, is not sufficient to show that it is a *partisan* gerrymander. To evaluate the partisan implications of the six plans, we take a precinct-level baseline voting pattern and aggregate it by district to explore hypothetical election outcomes under the six plans and the reference maps. The baseline pattern is calculated by averaging the vote totals for the three presidential elections and three gubernatorial elections that were held in Pennsylvania from 2000 to 2010.⁴ These election data were also used during litigation. While being far from a perfect way to create counterfactual election outcomes, this simple averaging of statewide results is often used in academic research and courts. The values of \hat{R} for the mean and various quantiles of the district-level vote shares were all within standard convergence ranges.

⁴Data from Ansolabehere, S. and Rodden, J. (2011). Pennsylvania Data Files. Available at <https://doi.org/10.7910/DVN/FJHHDS>.

Then, within each plan, we number the districts by their baseline Democratic two-party vote share, so District 1 is the least Democratic and District 18 the most. Figure 7(c), analogous to Figure 7 in [Herschlag, Ravier and Mattingly \(2017\)](#), presents the distribution of the Democratic two-party vote share for each of the districts across the reference maps and also shows the values for the General Assembly plan (triangles) and the court's adopted plan (circles). When compared to the reference maps and the court's plan, the General Assembly plan tends to yield smaller Democratic vote share in competitive districts (p -values of 0.0162, 0.0002, 0.0002, and 0.0002 for Districts 9–12 compared to the reference set) while giving larger Democratic vote share in noncompetitive districts. The effect of this is to produce six Democratic seats, on average, under the General Assembly plan, compared to a range of eight to 12 under the reference sample (p -value 0.0002) and 11 under the Court plan. All of these findings provide evidence that the General Assembly plan was gerrymandered in favor of Republicans by packing Democratic voters into noncompetitive districts.

7. Concluding remarks. Redistricting sampling algorithms allow for the empirical evaluation of a redistricting plan by generating alternative plans under a certain set of constraints. Researchers and policymakers can compute various statistics from the redistricting plan of interest and compare them with the corresponding statistics based on these sampled plans. Unfortunately, existing approaches often struggle when applied to real-world problems, owing to the scale of the problems and the number of constraints involved.

The SMC algorithm presented here can provably asymptotically sample from a specific target distribution. Because of the constructive sampling approach, it does not face the problem of potential reducibility that MCMC algorithms have and performs more efficiently in a realistic applied setting than a MCMC algorithm that uses an identical proposal distribution. The proposed method also incorporates, by design, the common redistricting constraints of population balance, geographic compactness, and minimizing administrative splits. We expect these advantages of the SMC algorithm to meaningfully improve the reliability of outlier analysis in real-world redistricting cases.

The proposed approach has some important limitations. Like other redistricting sampling algorithms, if the target distribution strays too far from the proposal distribution, the quality of the generated sample will suffer dramatically. Close examination of diagnostic measures, like \hat{R} , standard errors, sample diversity, and per-iteration efficiency measures, like weight variance and the number of surviving plans, are required to provide confidence that the algorithm is sampling accurately and has converged for particular quantities of interest. Our open-source software provides these diagnostics so that analysts can use the proposed algorithm successfully.

One important drawback particular to the SMC algorithm arises in situations with dozens or hundreds of separate districts. Like any other SMC or particle resampling scheme, as the number of iterations (districts) grows, eventually all of the samples will share a common ancestor. While this is not a problem in many SMC applications (such as Bayesian inference), for redistricting, this means that all of the sampled plans will share one or more districts that are completely identical. Any summary statistics which rely on these districts will have very large standard errors and may not converge, even with a large number of samples. While this problem will be readily identified with the diagnostics we propose here, there is little the analyst can do to address the issue beyond increasing the sample size and possibly combining many independent runs of the algorithm. One interesting avenue for future research would be to examine whether several applications of an MCMC kernel at various points in the SMC algorithm could help refresh the sample and counteract the tendency to collapse to a single ancestor.

Future research should also explore the possibility of improving several design choices in the algorithm to further increase its efficiency. Wilson's algorithm, for instance, can be generalized to sample from edge-weighted graphs. Choosing weights appropriately could lead to trees which induce maps that are more balanced or more compact. And the procedure for choosing edges to cut, while allowing for the sampling probability to be calculated, introduces inefficiencies by leading to many rejected maps. Further improvements in either of these areas should allow us to better sample and investigate redistricting plans over large maps and with even more complex sets of constraints.

Finally, we believe that the application of these sampling algorithms to real-world redistricting problems is essential. We have applied the SMC algorithm to the 2020 congressional redistricting in all 50 states (McCartan et al. (2022)), and this allowed us to evaluate the degree of partisan gerrymandering across states (Kenny et al. (2023)). Such applications shed light on substantive political science theories and highlight new methodological challenges.

Acknowledgments. The authors would like to thank Moon Duchin, Ben Fifield, Greg Herschlag, Mike Higgins, Chris Kenny, Jonathan Mattingly, Justin Solomon, and Alex Tarr for helpful comments and conversations. Imai thanks Yunkyu Sohn for his contributions at an initial phase of this project. We also thank the Editor, Associate Editor, and anonymous reviewers for detailed and helpful comments.

SUPPLEMENTARY MATERIAL

Appendices (DOI: [10.1214/23-AOAS1763SUPPA](https://doi.org/10.1214/23-AOAS1763SUPPA); .pdf). Proofs of propositions, additional validation example, and algorithm implementation details.

Replication materials (DOI: [10.1214/23-AOAS1763SUPPB](https://doi.org/10.1214/23-AOAS1763SUPPB); .zip). Replication code and data for the validation study, applied Pennsylvania analyses, and all appendices.

REFERENCES

- AKITAYA, H. A., KORMAN, M., KORTEN, O., SOUVAINE, D. L. and TÓTH, C. D. (2022). Reconfiguration of connected graph partitions via recombination. *Theoret. Comput. Sci.* **923** 13–26. [MR4436558 https://doi.org/10.1016/j.tcs.2022.04.049](https://doi.org/10.1016/j.tcs.2022.04.049)
- AUTRY, E., CARTER, D., HERSCHLAG, G., HUNTER, Z. and MATTINGLY, J. (2020). Multi-scale merge-split Markov chain Monte Carlo for redistricting. Working paper.
- BANGIA, S., GRAVES, C. V., HERSCHLAG, G., KANG, H. S., LUO, J., MATTINGLY, J. C. and RAVIER, R. (2017). Redistricting: Drawing the line. arXiv preprint [arXiv:1704.03360](https://arxiv.org/abs/1704.03360).
- BOZKAYA, B., ERKUT, E. and LAPORTE, G. (2003). A tabu search heuristic and adaptive memory procedure for political districting. *European J. Oper. Res.* **144** 12–26.
- CANNON, S., DUCHIN, M., RANDALL, D. and RULE, P. (2022). Spanning tree methods for sampling graph partitions. arXiv preprint [arXiv:2210.01401](https://arxiv.org/abs/2210.01401).
- CARTER, D., HERSCHLAG, G., HUNTER, Z. and MATTINGLY, J. (2019). A merge-split proposal for reversible Monte Carlo Markov Chain sampling of redistricting plans. arXiv preprint [arXiv:1911.01503](https://arxiv.org/abs/1911.01503).
- CHATTERJEE, S. and DIACONIS, P. (2018). The sample size required in importance sampling. *Ann. Appl. Probab.* **28** 1099–1135. [MR3784496 https://doi.org/10.1214/17-AAP1326](https://doi.org/10.1214/17-AAP1326)
- CHEN, J. (2017). Expert report of Jowei Chen, Ph.D. Expert witness report in League of Women Voters v. Commonwealth.
- CHEN, J. and RODDEN, J. (2013). Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Q. J. Polit. Sci.* **8** 239–269.
- CHIKINA, M., FRIEZE, A. and PEGDEN, W. (2017). Assessing significance in a Markov chain without mixing. *Proc. Natl. Acad. Sci. USA* **114** 2860–2864. [MR3628186 https://doi.org/10.1073/pnas.1617540114](https://doi.org/10.1073/pnas.1617540114)
- CHIKINA, M., FRIEZE, A. and PEGDEN, W. (2019). Understanding our Markov chain significance test: A reply to Cho and Rubinstein-Salzedo. *Stat. Public Policy* **6** 50–53.
- CHO, W. K. T. and LIU, Y. Y. (2018). Sampling from complicated and unknown distributions: Monte Carlo and Markov Chain Monte Carlo methods for redistricting. *Phys. A* **506** 170–178. [MR3810349 https://doi.org/10.1016/j.physa.2018.03.096](https://doi.org/10.1016/j.physa.2018.03.096)

- CHO, W. K. T. and RUBINSTEIN-SALZEDO, S. (2019). Understanding significance tests from a non-mixing Markov chain for partisan gerrymandering claims. *Stat. Public Policy* **6** 44–49.
- CIRINCIONE, C., DARLING, T. A. and O'ROURKE, T. G. (2000). Assessing South Carolina's 1990s congressional districting. *Polit. Geogr.* **19** 189–211.
- COVER, T. M. and THOMAS, J. A. (2006). *Elements of Information Theory*, 2nd ed. Wiley Interscience, Hoboken, NJ. MR2239987
- DEFORD, D., DUCHIN, M. and SOLOMON, J. (2021). Recombination: A family of Markov chains for redistricting. *Harv. Data Sci. Rev.* <https://doi.org/10.1162/99608f92.eb30390f>
- DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 411–436. MR2278333 <https://doi.org/10.1111/j.1467-9868.2006.00553.x>
- DOUCET, A., DE FREITAS, N. and GORDON, N., eds. (2001). *Sequential Monte Carlo Methods in Practice. Statistics for Engineering and Information Science*. Springer, New York. MR1847783 <https://doi.org/10.1007/978-1-4757-3437-9>
- DUBE, M. P. and CLARK, J. T. (2016). Beyond the circle: Measuring district compactness using graph theory. In *Annual Meeting of the Northeastern Political Science Association*.
- DUCHIN, M. (2018). Outlier analysis for Pennsylvania congressional redistricting.
- FIFIELD, B., HIGGINS, M., IMAI, K. and TARR, A. (2020a). Automated redistricting simulation using Markov chain Monte Carlo. *J. Comput. Graph. Statist.* **29** 715–728. MR4191238 <https://doi.org/10.1080/10618600.2020.1739532>
- FIFIELD, B., IMAI, K., KAWAHARA, J. and KENNY, C. T. (2020b). The essential role of empirical validation in legislative redistricting simulation. *Stat. Public Policy* **7** 52–68.
- GELMAN, A., CARLIN, J. B., STERN, H. S., DUNSON, D. B., VEHTARI, A. and RUBIN, D. B. (2014). *Bayesian Data Analysis*, 3rd ed. *Texts in Statistical Science Series*. CRC Press, Boca Raton, FL. MR3235677
- GELMAN, A. and RUBIN, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statist. Sci.* **7** 457–472.
- GUTH, L., NIEH, A. and WEIGHILL, T. (2022). Three applications of entropy to gerrymandering. In *Political Geometry—Rethinking Redistricting in the US with Math, Law, and Everything in Between* 275–292. Birkhäuser/Springer, Cham. MR4436942 https://doi.org/10.1007/978-3-319-69161-9_14
- HERSCHLAG, G., RAVIER, R. and MATTINGLY, J. C. (2017). Evaluating partisan gerrymandering in Wisconsin. arXiv preprint [arXiv:1709.01596](https://arxiv.org/abs/1709.01596).
- KENNY, C. T., MCCARTAN, C., FIFIELD, B. and IMAI, K. (2020). redist: Computational algorithms for redistricting simulation. <https://CRAN.R-project.org/package=redist>.
- KENNY, C. T., MCCARTAN, C., SIMKO, T., KURIWAKI, S. and IMAI, K. (2023). Widespread partisan gerrymandering mostly cancels nationally, but reduces electoral competition. *Proc. Natl. Acad. Sci. USA* **120** e2217322120.
- KOSTOCHKA, A. V. (1995). The number of spanning trees in graphs with a given degree sequence. *Random Structures Algorithms* **6** 269–274. MR1370962 <https://doi.org/10.1002/rsa.3240060214>
- LEE, A. and WHITELEY, N. (2018). Variance estimation in the particle filter. *Biometrika* **105** 609–625. MR3842888 <https://doi.org/10.1093/biomet/asy028>
- LEGLAND, F. and OUDJANE, N. (2005). A sequential particle algorithm that keeps the particle system alive. In *2005 13th European Signal Processing Conference* 1–4. IEEE.
- LIU, J. S., CHEN, R. and LOGVINENKO, T. (2001). A theoretical framework for sequential importance sampling with resampling. In *Sequential Monte Carlo Methods in Practice. Stat. Eng. Inf. Sci.* 225–246. Springer, New York. MR1847794
- LIU, J. S., CHEN, R. and WONG, W. H. (1998). Rejection control and sequential importance sampling. *J. Amer. Statist. Assoc.* **93** 1022–1031. MR1649197 <https://doi.org/10.2307/2669846>
- LIU, Y. Y., CHO, W. K. T. and WANG, S. (2016). PEAR: A massively parallel evolutionary computation approach for political redistricting optimization and analysis. *Swarm Evol. Comput.* **30** 78–92.
- MACMILLAN, W. (2001). Redistricting in a GIS environment: An optimisation algorithm using switching-points. *J. Geogr. Syst.* **3** 167–180.
- MAGLEBY, D. B. and MOESSON, D. B. (2018). A new approach for developing neutral redistricting plans. *Polit. Anal.* **26** 147–167.
- MATTINGLY, J. C. and VAUGHN, C. (2014). Redistricting and the will of the people. arXiv preprint [arXiv:1410.8796](https://arxiv.org/abs/1410.8796).
- MCCARTAN, C. and IMAI, K. (2023). Supplement to “Sequential Monte Carlo for sampling balanced and compact redistricting plans.” <https://doi.org/10.1214/23-AOAS1763SUPPA>, <https://doi.org/10.1214/23-AOAS1763SUPPB>
- MCCARTAN, C., KENNY, C. T., SIMKO, T., GARCIA III, G., WANG, K., WU, M., KURIWAKI, S. and IMAI, K. (2022). Simulated redistricting plans for the analysis and evaluation of redistricting in the United States. *Sci. Data* **9** 689.

- MCKAY, B. D. (1981). Spanning trees in random regular graphs. In *Proceedings of the Third Caribbean Conference on Combinatorics and Computing (Bridgetown, 1981)* 139–143. Univ. the West Indies, Cave Hill Campus, Barbados. MR0657198
- MEHROTRA, A., JOHNSON, E. L. and NEMHAUSER, G. L. (1998). An optimization based heuristic for political districting. *Manage. Sci.* **44** 1100–1114.
- NATIONAL CONFERENCE OF STATE LEGISLATURES (2021). Redistricting criteria. Available at <https://www.ncsl.org/research/redistricting/redistricting-criteria.aspx>.
- LEAGUE OF WOMEN VOTERS V. COMMONWEALTH (2018). 178 A. 3d 737 (Pa: Supreme Court).
- OLSSON, J. and DOUC, R. (2019). Numerically stable online estimation of variance in particle filters. *Bernoulli* **25** 1504–1535. MR3920380 <https://doi.org/10.3150/18-bej1028>
- PETERS, G. W., FAN, Y. and SISSON, S. A. (2012). On sequential Monte Carlo, partial rejection control and approximate Bayesian computation. *Stat. Comput.* **22** 1209–1222. MR2992295 <https://doi.org/10.1007/s11222-012-9315-y>
- POLSBY, D. D. and POPPER, R. D. (1991). The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale Law Policy Rev.* **9** 301–353.
- TUTTE, W. T. (1984). *Graph Theory. Encyclopedia of Mathematics and Its Applications* **21**. Addison-Wesley Company, Reading, MA. Advanced Book Program. With a foreword by C. St. J. A. Nash-Williams. MR0746795
- VEHTARI, A., GELMAN, A., SIMPSON, D., CARPENTER, B. and BÜRKNER, P.-C. (2021). Rank-normalization, folding, and localization: An improved \hat{R} for assessing convergence of MCMC (with discussion). *Bayesian Anal.* **16** 667–718. Includes comments and discussions by seven discussants and a rejoinder by the authors. MR4298989 <https://doi.org/10.1214/20-ba1221>
- WILSON, D. B. (1996). Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)* 296–303. ACM, New York. MR1427525 <https://doi.org/10.1145/237814.237880>
- WU, L. C., DOU, J. X., SLEATOR, D., FRIEZE, A. and MILLER, D. (2015). Impartial redistricting: A Markov Chain approach. arXiv preprint [arXiv:1510.03247](https://arxiv.org/abs/1510.03247).