Political Science Research and Methods (2026), 1–22
doi:10.1017/psrm.2025.10050

ORIGINAL ARTICLE

A Bayesian mixture model captures temporal and spatial structure of voting blocs within longitudinal referendum data

John O'Brien

Department of Mathematics, Bowdoin College, Brunswick, ME, USA
Email: jobrien@bowdoin.edu

(Received 17 October 2023; revised 14 January 2025; accepted 17 January 2025)

Abstract

The estimation of voting blocs is an important statistical inquiry in political science. However, the scope of these analyses is usually restricted to roll call data where individual votes are directly observed. Here, we examine a Bayesian mixture model with Dirichlet-multinomial components to infer voting blocs within longitudinal referendum data. This model infers voting bloc mixture within municipalities using state-level data aggregated at the municipal level. As a case study, we analyze the vote totals of Maine referendum questions balloted from 2008 to 2019 for 423 municipalities. Using a birth–death Markov chain Monte Carlo approach to inference, we recover the posterior distribution on the number of voting blocs, the support for each question within each bloc, and the blocs' mixture within each municipality. We find that these voting blocs are structured by geography and largely consistent across the study period. The model finds that blocs exhibit both spatial gradients and discontinuities in their structure. Examining the statistical fit of the model, we uncover a small number of questions that show inconsistency with the statewide bloc structure and note that the content of these questions relates to specific regions. We conclude with possible statistical extensions, connections to other statistical frameworks in political science, and detail possible locations for model applications.

Keywords: Bayesian; latent class models; quantitative methods; voting behavior

1. Introduction

1.1. Statistical models of voting blocs

An electorate of more than a few individuals is assured to be a mixture of different political opinions. What factors explain the structure and persistence underlying this heterogeneity of dispositions and how they connect to voting behavior is a fundamental question of political behavior, both within political science and across the social sciences (Gallup and Rae, 1940; May, 1973; Converse, 1987; Weidlich, 1994; Gunn, 1995; Kinder, 1998; Weakliem and Biggert, 1999; Leeper and Slothuus, 2014). Consequently, there is a vast statistical literature to ascertain the association of possible explanatory variables associate with vote totals. This overall logic—proceeding from explanatory variables to explain aspects of aggregated voting behavior—finds expression in commonly used statistical frameworks in political science such as generalized regression, ecological regression, and latent class analysis (McCutcheon, 1987; King, 1988; Beck and Jackman, 1998; Schuessler, 1999; King et al., 2004; Magidson et al., 2020) each the site of rich statistical innovation and discussion (King, 1990; Beck, 2000; Gelman et al., 2001; Lanza et al., 2013).

© The Author(s), 2026. Published by Cambridge University Press on behalf of EPS Academic Ltd. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

Here we take up an alternative approach to understanding the structures underlying aggregated vote totals: we seek to identify self-consistent patterns of association—voting blocs—within voting data itself using a Bayesian mixture model (McLachlan and Basford, 1988; Frühwirth-Schnatter, 2006; Fruhwirth-Schnatter et al., 2019; McLachlan et al., 2019). In this regard, this approach is more closely aligned with models in genetics, metagenomics, and topic modeling, where mixture modeling is used as an empirical complement to theoretical explanatory frameworks (Blei et al., 2003; Falush et al., 2003; Holmes et al., 2012; Hellenthal et al., 2014). While the model-based approach here is novel, the broader interest in constructing voting blocs from aggregated vote totals within political science is long-standing (Snyder Jr, 1996). The approach presented here provides several extensions to existing voting bloc methods that have largely focused on vote data from single individuals in the context of parliamentary or deliberative voting: (1) it shows how to infer voting blocs from aggregated vote data rather than individual votes, allowing extension of voting bloc analysis to a common form of election data—referendum vote totals; (2) it introduces a novel inferential method into the political science literature—birth–death Markov chain Monte Carlo (BD-MCMC)—for determining the number of voting blocs; and, (3) provides an novel method to ascertain the statistical consistency of responses to referendum questions. The model also permits an examination of the spatial and temporal structure of these voting blocs, which may permit associations with structural features, such as migration history, demographics, or economic modalities. While not considered here, this approach can also point toward new methods for polling that take into account long-term latent structures. We discuss each of these points in greater depth in theDiscussion section.

In the context of political science, our statistical approach follows in the footsteps of Gormley and Murphy (2008a, 2008b), who developed a set of mixture models to identify voting blocs from rank choice voting (RCV) election data(Gormley and Murphy, 2006, 2008a, 2008b; Gormley and Frühwirth-Schnatter, 2019). With these data, Gormley and Murphy use an approach that leverages the covariance structure within these preferences to infer four voting blocs within the Irish presidential electorate, each blocs' voting preferences, and their overall proportions within the population. In a later work, they show how external covariates can be associated with the voting blocs using mixture-of-experts models (Gormley and Murphy, 2010). Unfortunately, the broad application of this approach in political science may have been undercut by the relative infrequency of RCV elections and consequently the limited availability of this type of data. The statistical goal of this paper is to extend their approach in two ways: (1) to adapt mixture models to the context of referendum vote totals, a more common form of election data, and (2) to capture not previously considered spatial and temporal variation within these data. To explore the ability of this model to explicate the structure of voter behavior we undertake a case study of municipal vote totals from 54 referendum questions in the US state Maine balloted between 2008 and 2019.

To uncover the structure of voting blocs from referendum vote totals aggregated at the municipal level, we use a Bayesian mixture model. Following Gormley and Murphy, we refer to voting blocs to describe voters with similar voting patterns. A voting pattern is constituted by a vector of distributions of support for each question. In constructing the model, we make the following assumptions about the data. We assume that there are a finite but unknown number of voting blocs shared by all municipalities within the state, that the municipal voting data is compositional with fixed mixture proportions, and that the voting patterns fully characterize the observed vote totals within each municipality(Aitchison, 1982). We further assume that voting blocs are mixed independently within each municipality and that each voting pattern is defined by a Dirichlet-multinomial (DM) distribution describing the mean and dispersion of support for each question. The assumption of independent mixture allows the data to "speak for itself" in determining each municipality's mixture proportion but comes at the cost of losing any inferential power that might be garnered by sharing information across neighboring communities. In the case study here, we restrict attention to just yes/no totals so the Dirichlet-multinomial distribution becomes a Beta-binomial (BB) distribution. Lastly, we assume the BB distributions are independent across questions within voting patterns. While data

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

across questions likely possess correlation structure, the assumption of independence provides the model maximum flexibility to determine these components from the data. This again comes at the cost of increased variance in the estimates that could be recovered with more involved prior structure, a point that we return to in the Discussion. The observed support for each question within each municipality is then derived by integrating the distribution of voting patterns over the mixture proportion of the voting blocs in each municipality. The inferential objective is then to recover the distribution of the number of voting blocs, the structure of their voting patterns (i.e. the distribution of support for each question), and the mixture proportions for each voting bloc within each municipality.

Mixture models, often known as model-based clustering approaches, constitute an important extension of many common dimensional reduction techniques such as principal components analysis (PCA) already broadly used in political science (Simon and Xenos, 2004; Pennings et al., 2005; Magyar, 2022). Model-based clustering, which has also made recent if less certain inroads in political science, provides an explicit probabilistic description of the data observation process, whereas more classical methods (like PCA) often rely on the construction of sensible metrics to operate effectively (Hill and Kriesi, 2001; Ahlquist and Breunig, 2012). Model-based clustering approaches can then be used to rigorize and extend more direct dimensional reduction methods to provide a richer understanding of aspects of the data. For instance, PCA can provide an incisive sense of the latent structures of the data but may be biased by choice of metric or missingness. It is also constrained to only allow the researcher to make "by eye" estimates of any clusters. Model-based clustering formalizes much of this process by the use of flexible probabilistic models to account for these issues at the cost of a significantly more complex and computationally intensive inference scheme. In general, researchers will often turn to PCA and similar techniques to first understand how clustered the data are and what other issues might be relevant for analysis before turning to a model-based clustering approach that requires significantly more mathematical and computational overhead. As we will see in the Results below, the mixture model approach here is well-complemented by PCA, showing overall consistency while also permitting analysis of more finely-grained clusters within the data.

While models for inferring latent structures like voting blocs directly from aggregated vote totals are uncommon, there is an extensive literature on estimating voting blocs—and, more generally, latent space structures—from roll call and similar data where the vote of each participant is observed, such as parliaments, judicial panels, and corporate boards (Katz and King, 1999; Clinton et al., 2004; Hix et al., 2006; Bailey et al., 2017; D'Angelo et al., 2019; Reuning et al., 2019). For a careful review of the models used in roll call voting, see McAlister (2020). These studies also frequently make use of elements of the mixture model employed here, though many other methods abound, including PCA, hierarchical clustering, latent class analysis, and topological data analysis (Holloway, 1990; De Leeuw, 2006; Amelio and Pizzuti, 2012; Vejdemo-Johansson et al., 2012; Bakk et al., 2014; Grimmer et al., 2022). The key distinction between these studies and the current work is that they analyze contexts where votes are directly observed at the individual level whereas in referendum election data, vote totals are necessarily aggregated at some governmental level (e.g. here, municipality). The use of latent space models and their relatives to understand referendum data has not been entirely neglected: two recent efforts to understand Swiss referendum results made use of network-based approaches and latent class models, respectively (Koseki, 2018; Mantegazzi, 2021) while an early unpublished paper on California referendum results executed an analysis with elements similar to both Gormley and Murphy's approach and the approach we develop here (Dubin and Gerber, 1992).

Determining the overall number of voting blocs is naturally a point of interest, both in its own right and as a means to integrate over uncertainty in the number of voting blocs. As the number of parameters in the model varies with the number of voting blocs, we require an approach to transdimensional Markov chain Monte Carlo (MCMC) integration to infer the full posterior distribution of the number of voting blocs (Green, 2003). Methods for transdimensional inference, while not yet widely used in political science, have been taken up to deal a number of political science problems

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

where uncertainty in the number of components is important, including estimating conflict regime changes during the American occupation of Iraq, marking changes in voter or donor behavior, and modeling changes in approval ratings (Spirling, 2007; Ratkovic and Eng, 2010; Blackwell, 2018; Kim, 2020; Park and Yamauchi, 2023). These methods provide additional statistical clarity over more frequently used criterion-based methods by furnishing not just estimates of the best or most probable model but an assessment of the relative probability over the set of relevant plausible models, permitting more finely grained comparison between model hypotheses. As Bayesian methods become more extensively utilized within political science, these approaches will also be increasingly required to fully infer complex model posteriors.

Here, we implement a version of the BD-MCMC algorithm to estimate the posterior distribution over the number of mixture components as well as their parameters, which we believe to be a novel contribution to the political science literature. A birth–death process is a continuous-time Markov chain that takes natural numbers as its values, jumping between neighboring values. The waiting times between transitions are distributed exponentially while transitions between states are of two types: births, where a state is added to the process, and deaths, where a state is removed from the process. Stephens (2000); Shi et al. (2002) showed how this general process could be encoded as an MCMC algorithm for mixture models to generate samples from the posterior distribution of the number of components. The specific focus of their algorithm was to estimate the number of clusters, though they showed how this could be extended to other contexts, such as change-point detection. This Markov chain operates similarly to the general process but with mixture models as the states, alternating between exponentially-distributed wait times and simulating birth–death moves where either a new mixture component is added to the model (a birth) or a component is deleted (a death). Their key insight was how to calibrate the wait time and the probabilistic selection of birth–death moves in order to satisfy a detailed balance condition that ensures that the amount of time spent at each number of components in the Markov chain converges to the posterior distribution on the number of components. While BD-MCMC approaches have become increasingly common in applied statistics for finding changepoints, estimating correlation graph structures, and determining the number of components, they have not yet found purchase in political science {Xu et al.,2022, Wang et al., 2020, Mohammadi et al., 2017). In particular, in medical imaging, a similar approach to the one presented here has been developed to determine the number of components underlying clinical data (Amirkhani et al., 2021).

1.2. Maine referendum data

Referendums are among the most direct forms of electoral democratic engagement, requiring a plebiscite of eligible voters on a specific question of governance, ranging from bonds to constitutional amendments(Bowler et al., 1998; Mendelsohn and Parkin, 2001), and form increasingly common way for governments to decide on issues of public controversy. Referendum elections are employed at the national level in at least 20 countries, 26 US states, and hundreds of cities[1] The two most common forms of referendum—the citizen initiative, where citizens can circumvent legislatures and directly query the voting populace, and popular referendum, where the citizenry either approves or refuses a legislative act—will be treated as the same for the purposes of this manuscript, although the frequency of the former vastly outweighs the latter in the data considered here (51 out of 54 questions). We restrict attention solely to yes and no votes as blank ballots were variably recorded in the raw data (Maine Bureau of Corporations, Elections, and Commissions, 2021).

The November referendum data from 2008 to 2019 from the US state Maine possesses several desirable qualities from a modeling perspective: a relatively large number of balloted questions; consistency in municipal boundaries and voting procedures over the study period; and a largely stable

[1]These numbers were derived using Wikipedia and Ballotpedia. We find no academic source that provides a more comprehensive examination.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

demography. Maine was the first state in the eastern US to adopt citizen-initiated referendum and has a vibrant tradition of these elections(Scontras, 2016). Since 1980, there have been more than 200 referendum questions and no year without a November ballot question until 2020 when Covid-19-related pandemic restrictions inhibited election procedures. Municipally, Maine is constituted in the New England town model with the municipality as the primary structure of local governance, with negligible population in either the compact populated places or unorganized territories common elsewhere in the US. Nearly all towns in the state were established by the beginning of the 20^{th} century, with only one town appearing during the study period (an island seceding from its mainland community). The Maine Bureau of Elections has kept largely consistent voting records throughout the study period, while similar quality analog records go back to the late 19^{th} century (Archives, 2022). We downloaded the referendum data for each November election in the study set from the Maine State Bureau of Corporations, Elections, and Commissions website on June 1, 2020 (Maine Bureau of Corporations, Elections, and Commissions, 2021). These raw files were processed according to the procedure outlined below and the cleaned files positioned on the Github page for this project: https://github.com/cascobayesian/Elections.BBMM.

Maine contains 721 municipalities, graded between cities, towns, plantations[2], and townships, in addition to a set of unorganized territories with negligible population. The number of municipalities present in each dataset varied, largely according to whether townships or plantations were aggregated into neighboring municipalities for a particular election. To homogenize the number of municipalities across years, we combine any municipality together with a neighboring municipality if the two were combined for some election in the study period. We also eliminate vote totals from unorganized territories as well as townships or plantations where the presence of a polling station was not consistent across the data set. Any remaining municipality that registered no votes for a question was considered an error and eliminated from the analysis. Through this elimination, we arrive at a list of 423 municipalities for the data set here.

Referendums were most frequently balloted at November elections, when congressional, gubernatorial, and presidential contests were also held, but were also administered during party primary elections in June. During the study period from 2008 to 2019, 61 questions were balloted in total with seven in June and 54 in November. Historically, June elections have both low and highly variable turnout relative to November elections, so we restrict attention to November questions. Table 1 provides additional information about the questions included in the study.

1.3. Outline of paper

The focus of the paper is to develop a Bayesian mixture model to infer voting blocs across twelve years of municipal referendum data and then interpret those results, with a focus on the statistical performance of the model. We first provide some notation and proceed to a generative description of the model. For inference, we describe an MCMC approach that employs a double replacement scheme to efficiently sample parameters; we also develop a BD-MCMC to infer the posterior distribution of the number of voting blocs and their parameters. As a complement to the case study, we perform a simulation to understand the inferential requirements of the algorithm in terms of number of municipalities, number of questions, and vote totals.

We present the results of the case study, first for all the questions analyzed together and then broken down by election cycle. We find that voting blocs are structured geographically with strong consistency across election cycles and that the voting blocs are organized both independently and along gradients. As a consistency check, we show that these voting blocs are consistent with information-theoretic projections of the data. Using posterior checks on the model, we examine the performance of the model by municipality and by question and uncover the presence of a small number of "local"

[2]The designation "plantation" is a distinct historical usage unique to Maine and denotes a unit of municipal self-government with less power than a town but more than a township.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

Table 1. Summary of referendums on each ballot in the study period. Total votes are measured as the highest total number of votes on each ballot. Median total refers to the median total number of votes for municipalities in the data set

Year	Num.	Total votes	Median total	Content
2008	3	707,915	805	Healthcare tax, casino, water bond
2009	7	562,985	673	Same-sex marriage, car efficiency rebate, school redistricting, government spending medical marijuana, transport bond, ballot reform
2010	3	554,783	639	Casino, dental care bond, conservation bond
2011	4	387,881	440	Same-day voter registration, casino, slot machine facility, redistricting change
2012	5	697,423	788	Same-sex marriage, education bond, sewer bond, transportation bond, water bond
2013	5	214,507	213	National Guard bond, two university bonds, transportation bond, community college bond
2014	7	597,224	694	Bear hunting, agriculture bond, small business bond, science bond, water bond, business incubator bond
2015	3	216,248	230	Clean elections, housing bond, transportation bond
2016	6	748,337	848	Marijuana legalization, income tax increase, gun background checks, minimum wage
2017	4	341,139	363	Casino, Medicaid expansion, transportation bond, government financial reform
2018	5	626,470	715	Home care program, sewer bond, transportation bond, university bond, community college bond
2019	2	186,207	167	Transportation bond, disabled signatures

questions that do not conform to the model's assumptions. We conclude with discussion of a number of further developments to the model, connections to other statistical approaches for understanding voting data, and other applicable locations/data sets.

2. Data and model

In the next two subsections, we describe the notation for the data and the statistical model used to analyze them. These subsections contain most of the notation used subsequently in the manuscript. For clarity, we summarize the notation in Table 2.

2.1. Data notation

The vote tables (yes/no) for the 423 consistent municipalities comprise the complete data set used in this study (Table 1). This complete data set is also broken down into 4-year election cycles, starting with 2008–2011 and ending with 2016–2019. We enumerate towns by i with $i = 1, \cdots, N = 423$. We index questions q from $q = 1, \cdots, 54$ and let $s \in \{0, 1\}$ denote if we refer to a "no" count or a "yes" count. For each town i, the observed counts are denoted:

$$\mathbf{c}_{iq} = (c_{iq0}, c_{iq1}),$$

where $c_{\cdot\cdot 1}$ indicates the "yes" votes and $c_{\cdot}$ indicates the "no" votes for the ballot question corresponding to q. We use $\boldsymbol{c}$ to denote the aggregation of all data; then we further denote $\boldsymbol{c}_{all}$ for the complete data and $\boldsymbol{c}_t$ for the four-year cycle starting with that year (e.g. $\boldsymbol{c}_{2008}$) for each of the four-year cycles.

2.2. Beta-binomial mixture model

We now lay out the Bayesian model for the data, starting with the data likelihood. As discussed above, we employ a finite mixture model to capture the voting blocs within the municipal referendum data. The model assumes that there are $k = 1, \cdots, K$ voting blocs defined by a distribution of support for all questions. For the purposes of initial exposition, we hold that K is fixed and then expand the model to treat it as a random variable. The mixture proportions for each municipality are specified by λ_{ik}

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

Table 2. Notation and interpretations for parameters of the statistical model of the data

Notation	Interpretation
Q	Number of referendum questions
N	Number of municipalities
K	Number of voting blocs (VBs)
$q = 1, \cdots, Q$	Index over referendum questions
$i = 1, \cdots, N$	Index over municipalities
$k = 1, \cdots, K$	Index over voting blocs
$\mathbf{c}_{iq} = (c_{iq0}, c_{iq1})$	No and yes vote totals for municipality i and question q
$\mathbf{c} = [\mathbf{c}_{iq} : i = 1, \cdots, N; q = 1, \cdots, Q]$	Data for all municipalities and questions
$\boldsymbol{\alpha}_{kq} = (\alpha_{kq0}, \alpha_{kq1})$	Beta-binomial parameters for VP k and question q
$\boldsymbol{\alpha} = [\boldsymbol{\alpha}_{kq} : k = 1, \cdots, K; q = 1, \cdots, Q]$	All Beta-binomial parameters
$\boldsymbol{z} = [z_1, \cdots, z_N],\ z_i \in \{1, \cdots, K\}$	Auxiliary variables assigning municipalities to voting blocs
$\boldsymbol{\lambda} = [\lambda_{ik}]$	Mixture proportion of voting bloc k in municipality i
η_k	Prior probability of municipality being in VP k
$\kappa,\ \theta$	Prior parameters for each $\boldsymbol{\alpha}_{kq}$

and subject to the natural restriction $\sum_{k=1}^{K} \lambda_{ik} = 1$. The inferential modeling goal is then to recover the number of voting blocs (K), the parameters for each of their component voting patterns, and the mixture proportions within each municipality.

Within each of the k voting blocs and for each question q, we use a BB distribution specified by parameters $\boldsymbol{\alpha}_{kq} = (\boldsymbol{\alpha}_{kq0}, \boldsymbol{\alpha}_{kq1})$ to describe the distribution for c_{iq}. The BB distribution generalizes the more familiar binomial distribution by integrating the support proportion p over a Beta distribution, parameterized by α_{kq0} and α_{kq1}, leading to a probability for the observed counts as:

$$\begin{aligned}
\pi(\boldsymbol{c}_{iq}|\boldsymbol{\alpha}_{kq}, z_i = k) &= \int_0^1 \text{BINOMIAL}(c_{iq0}|c_{iq0} + c_{iq1}, p) \cdot \text{BETA}(p|\alpha_{kq0}, \alpha_{kq1}) dp \\
&= \frac{1}{B(\alpha_{kq0}, \alpha_{kq1})} \int_0^1 \binom{c_{iq0} + c_{iq1}}{c_{iq0}} p^{c_{iq0}} \cdot (1-p)^{c_{iq1}} \cdot p^{\alpha_{kq0}-1} \cdot (1-p)^{\alpha_{kq1}-1} dp \\
&= \binom{c_{iq0} + c_{iq1}}{c_{iq0}} \frac{B(c_{ik0} + \alpha_{kq0}, c_{ik0} + \alpha_{kq1})}{B(\alpha_{kq0}, \alpha_{kq1})},
\end{aligned}$$

where $B(\alpha, \beta)$ is the Beta function and $BETA(\alpha, \beta)$ is the Beta distribution. This distribution has the advantage of allowing flexibility in capturing both the mean support and its dispersion, as well as permitting limited bimodality.

As there are K voting blocs, the probability of the data for municipality i and question q is then the sum of BB densities for each k weighted by λ_{ik} :

$$\pi(\mathbf{c}_{iq}|\boldsymbol{\alpha}_{\cdot q}) = \sum_{k=1}^{K} \lambda_{ik} \cdot \pi(\mathbf{c}_{iq}|\boldsymbol{\alpha}_{kq}). \tag{1}$$

The summation within this probability creates a significant challenge for MCMC inference as individual component parameters cannot be easily updated (McLachlan et al., 2019). We therefore make use of latent variable data augmentation, a standard technique in Bayesian mixture inference to recast the model to a more computationally accessible form (Frühwirth-Schnatter, 2006; Fruhwirth-Schnatter et al., 2019). To do this, we augment the data with latent variables $z_i \in \{1, \cdots, K\}$ for $i = 1, \cdots, N$ that specify the voting bloc for each municipality i for all questions. This formulation also requires that we specify the prior probabilities of a municipality being part of component k, which we label $\boldsymbol{\eta} = (\eta_1, \cdots, \eta_K)$ with the natural restriction that $\sum_{k=1}^{K} \eta_k = 1$. The mixture parameters λ_{ik} can then be recovered from the posterior distribution, as detailed in the Posterior Inference subsection below.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

The latent variable data augmentation formulation then allows the probability of the data for municipality i and question q to be written as:

$$\begin{aligned}\pi(\mathbf{c}_{iq}|\boldsymbol{\alpha}, z_i, \boldsymbol{\eta}) &= \pi(z_i = k|\boldsymbol{\eta}) \cdot \pi(\mathbf{c}_{iq}|\boldsymbol{\alpha}_{z_i q}) \\ &= \eta_{z_i} \pi(\mathbf{c}_{iq}|\boldsymbol{\alpha}_{z_i q \cdot}).\end{aligned}$$

Knowing the position of each municipality in each bloc removes the summation in Equation 1. Conditional on z_i and the $\boldsymbol{\alpha}$ parameters, each municipality's vote probability is then given by a single BB distribution. This allows the augmented data likelihood to be a product over all questions and municipalities:

$$\begin{aligned}\pi(\mathbf{c}|\boldsymbol{\alpha}, \boldsymbol{z}, \boldsymbol{\eta}) &= \pi(\mathbf{c}|\boldsymbol{\alpha}, \boldsymbol{z}) \cdot \pi(\boldsymbol{z}|\boldsymbol{\eta}) = \prod_{i=1}^{N} \pi(z_i|\boldsymbol{\eta}) \cdot \pi(\mathbf{c}_{iq}|\boldsymbol{\alpha}_{z_i q}, z_i) \\ &= \prod_{i=1}^{N} \eta_{z_i} \left\{ \prod_{q=1}^{Q} \pi(\mathbf{c}_{iq}|\boldsymbol{\alpha}_{z_i q}) \right\} \\ &= \prod_{i=1}^{N} \eta_{z_i} \left\{ \prod_{q=1}^{Q} \binom{c_{iq0} + c_{iq1}}{c_{iq0}} \frac{B(c_{iz_i 0} + \alpha_{kq0}, c_{iz_i 1} + \alpha_{z_i q 1})}{B(\alpha_{z_i q 0}, \alpha_{z_i q 1})} \right\}. \end{aligned} \tag{2}$$

This likelihood can be understood either as a finite mixture model where questions are treated as independently (but not identically) distributed data replicates or as a degenerate hidden Markov model with a fixed class across years. We return to these two perspectives in the discussion when we take up possible model extensions.

2.3. Prior specifications and generative model

Completion of the Bayesian model requires specifications for the priors for each of the parameters. As $\boldsymbol{\alpha}$ is the aggregation of $\boldsymbol{\alpha}_{kq} = (\alpha_{kq0}, \alpha_{kq1})$, for $k = 1, \cdots, K$ and $q = 1, \cdots, Q$, we assume that each of these parameters is independently drawn from a Gamma distribution with shape parameter κ and scale parameter θ. This choice provides a convenient Gibbs update discussed below. Since κ and θ are the same for α_{kq0} and α_{kq1}, this amounts to a weakly informative prior on the proportion of support for a question to be centered at $\frac{1}{2}$. For all runs here, $\kappa = 1$ and $\theta = 10$. Numerical experiments indicate that even small municipalities' vote totals are sufficient to correctly infer the parameter values.

As noted above, the prior probability of z_i taking value k is η_k, completing the prior specification. The vector $\boldsymbol{\eta} = (\eta_1, \cdots, \eta_K)$ is a set of proportions and so sums to one. Absent any information about the relative weights of the components, a symmetric Dirichlet parameterized by $\mathbf{1}_K$, a vector of K ones, provides a reasonable, weakly informative prior. We complete the model's prior specification by assuming that K, the number of voting patterns, is distributed as a Poisson random variable with mean λ. For all runs considered here, we set $\lambda = 10$. The model can be compactly summarized as a generative procedure:

$$\begin{aligned} K &\sim \text{POISSON}(\lambda) \\ \boldsymbol{\eta} \mid K &\sim \text{DIRICHLET}(\mathbf{1}_K) \\ z_i \mid \boldsymbol{\eta} &\sim \text{CATEGORICAL}(\boldsymbol{\eta}) \\ \alpha_{kqs} &\sim \text{GAMMA}(\kappa, \theta) \\ \boldsymbol{c}_{iq}|\boldsymbol{\alpha}, z_i &\sim \text{BETA-BINOMIAL}(\alpha_{z_i q 0}, \alpha_{z_i q 1}), \end{aligned}$$

where $\mathbf{1}_K$ denotes a vector of K ones.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

2.4. Model inference preliminaries

The model presented has a Bayesian structure, and so the inferential goal is to recover the posterior distribution of the model parameters. As analytic approaches are not tractable, we turn to MCMC, a numerical technique for drawing samples from the posterior distribution. Of particular interest in this sampling process is K, the number of mixture components, $\boldsymbol{\alpha}$, the BB parameters specifying question support for each voting bloc, and $\boldsymbol{z} = [z_i : i = 1, \cdots, N]$, the mixture proportions for each voting bloc in each community. The mixture parameters $\boldsymbol{\lambda}$ are derived from the posterior distribution on $\boldsymbol{z}$ and so discussed separately from the presentation of the MCMC algorithms. As K changes the number of parameters in $\boldsymbol{\alpha}$ also changes, so the problem is transdimensional in nature. Consequently, we require an appropriate inference approach to for this context. For this, we employ a BD-MCMC algorithm. As the BD-MCMC algorithm is novel in the political science literature, we separate the presentation of that section of the algorithm from that for the finite-dimensional MCMC algorithm with fixed K. We outline the MCMC algorithm with fixed K and then describe the birth–death process leading to the BD-MCMC. We provide an Appendix for the complete details of the algorithm.

3. Results

We apply the model to the data set in two complementary ways: first to all questions aggregated into a single data set ($Q = 54$); and second, with separate data from each of nine four-year election cycles (e.g. 2008–2011). For these subsets, Q varies from sixteen to twenty-one.

3.1. Complete analysis indicates voting blocs' geographical structure

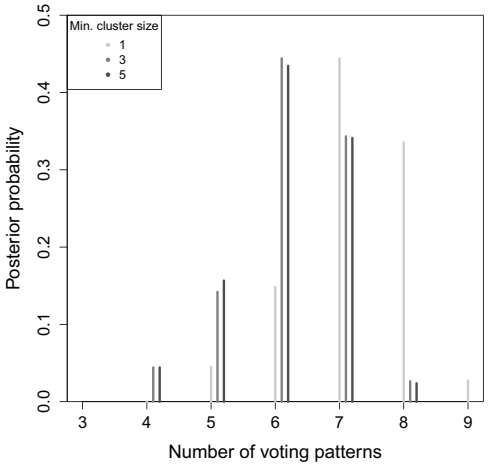
Figure 1 summarizes the results of the model applied to the complete data set, showing the posterior distribution on K, the confusion matrix across the posterior samples, and projection of the confusion matrix onto a map of the state. The posterior distribution for each K (Figure 1(a)) is estimated as a weighted mean of the BD-MCMC samples' K values weighted by their jump wait times. Observing that most posterior samples possess transient voting bloc clusters, we filter estimates of K to only count components that include a minimum number of members (1, 3, and 5). This filtering generates a narrowed and converging range of posterior values as the minimum number increases. The mode of the distribution is at $K = 6$, with significant mass at $K = 7$ and $K = 5$. We observe little variability across the runs in the distribution of K, though two runs slightly preferred the posterior mode at 7 rather than 6. In presenting the graphical results and the mixture proportions, we select $K = 6$ as a representative value though similar presentations are possible for $K = 5$ and $K = 7$.

In Figure 1(b), we summarize the transdimensional cluster structure using a matrix of pairwise co-occupancy frequency for each pair of municipalities, weighted according to the wait times in the BD-MCMC. We then sort both the rows and columns of the co-occupancy matrix according to a k-medoid clustering applied to the raw matrix with $K = 6$. This procedure has two significant advantages over alternatives: it is insensitive to the label switching problem since all calculations are only on pairwise municipal membership; and it effectively articulates the overall cluster structure while taking a single K as representative. We refer to this as the *representative clustering*.

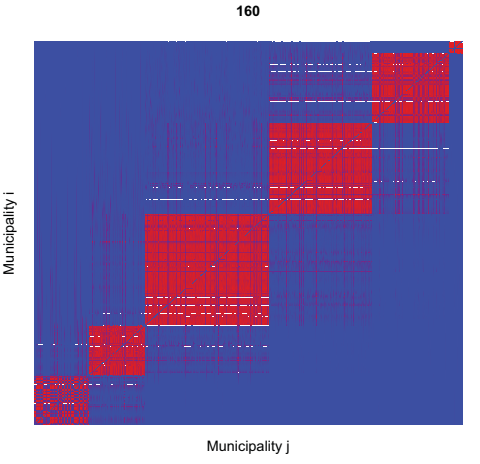
To inspect possible spatial structure within the voting bloc distribution, we project the results of the co-occupancy matrix onto a map of Maine, yielding Figure 1(c). For each municipality, we take the median occupancy found in the co-occupancy matrix and then normalize it to find the mixture proportions on the map. We find this statistic to be significantly more robust than direct estimates of bloc membership from the model that are again burdened by both label switching and the variability in K. We color the resulting voting blocs from south to north, with the southernmost municipality with majority occupancy in that voting bloc being used to determine position.

To provide a more accessible understanding of the voting blocs, we name them according to their geography and economic character. The first voting bloc (in red) is formed by many of the oldest

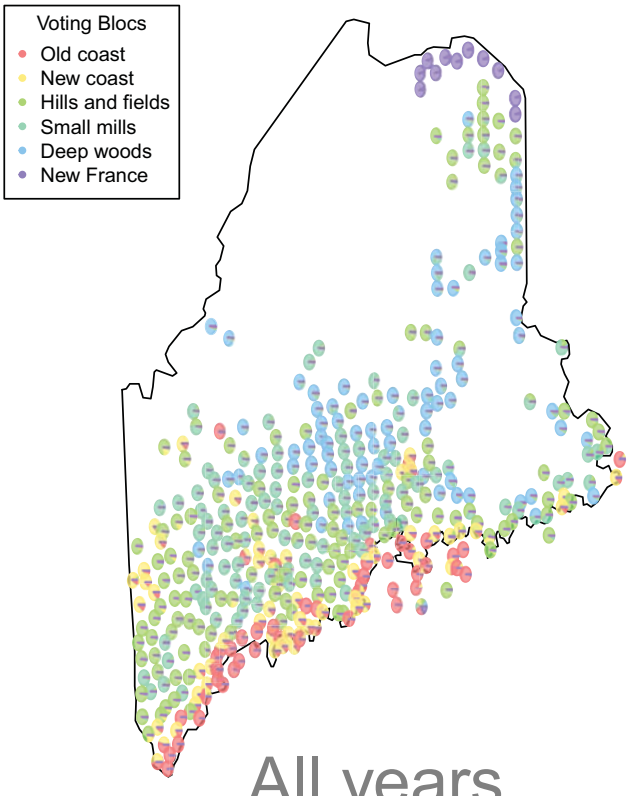
https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press



(a)



(b)



(c)

Figure 1. (a) Posterior distribution for K for complete data set. As small clusters are often transient, the results after filtering these clusters are also presented. (b) Co-occupancy matrix for the voting blocs of the complete data, using the representative clustering for $K = 6$. Color scale is blue at 0 and red at 0.96. (c) Projection of the voting bloc co-occupancy matrix onto the map of Maine. Voting blocs names are described in the Results section. Cluster (voting bloc) proportions calculated against the representative clustering.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

communities, largely found on the coast, we label as The Old Coast, and are distinguished by strong support for nearly all bonds, cultural changes (e.g. gay marriage), opposition to casinos, health care reform, and additional taxes on the wealth. The second voting bloc (in yellow) constitutes The New Coast, settled later in the colonial period but largely along rivers or their coastal outlets and distinguished electorally by similar positions to the Old Coast but with less support for gay marriage, gun regulation, hunting regulation, or health care expansion. The third voting bloc (in light green) we call the Hills and Fields, as these are largely agricultural communities that were truck farms through most of the 20^{th} century. They exhibit less support for gay marriage, gun regulation, and health care expansions than the New Coast and more support for casinos. The fourth voting bloc (in blue-green) we term the Small Mills, as these are small communities, largely lying tributary rivers, distinguished by a long history of light manufacturing and resource extraction. Electorally, they show even less support for gay marriage, gun regulation, and health care expansion than the Hills and Fields as well as reduced support for most bonds. The fifth voting bloc we call the Deep Woods, as these are the most rural communities are all largely organized around forestry. They exhibit even less support for gun regulation and gay marriage, as well as markedly low support for transportation bonds, education bonds, and health care expansion. Following from the book *American Nations*, the final bloc we name New France, as it constitutes the French-speaking communities (Woodard, 2012). Electorally, these communities are most similar to the New Coast but with much higher support for casinos, health care expansion, and gay marriage, and much lower support for marijuana legalization.

3.2. Consistency of voting blocs across nine 4-year election cycles

To explore the temporal patterns within the data, we apply the model to the data separated into nine four-year election cycles. The number of ballot questions in each cycle varies from 16 to 21 with a median of 18. As noted in the simulation study in the Appendix, this falls slightly below the number of questions required for consistent inference so additional variation in the results is to be expected. However, as the simulation was conducted with more highly mixed samples than observed in the data, Figure 2 shows the representative clustering for each election cycle.

No cycle shows results identical to the complete data set, though most sustain the key patterns found there: a set of coastal voting blocs also including inland cities; a set of voting blocs covering the inland municipalities excluding cities; and a distinct voting bloc in the far north covering the French-speaking communities. The primary difference from the complete data results lies in the number of blocs that cover the coastal and the inland municipalities. For instance, in the 2016 cycle, the southern coastal municipalities and the inland cities are covered by three voting blocs instead of two voting blocs as in the complete data.

There are two significant departures in the model results between the election cycle data and the complete data that appear to arise from the inclusion or exclusion of specific questions. The first departure is the emergence of a distinct bloc in the 2008 and 2009 cycles, covering the central, eastern, and English-speaking northern municipalities (bloc 5 in the 2008 cycle, bloc 6 in the 2009 cycle). This localizes the inland bloc more substantially than in the complete data model. We note that these two cycles contain a high number of casino-related questions (4 of 16 questions in the 2008 cycle, 3 of 18 in the 2009 cycle) whose unusual properties are explored below. The second departure is the absence of a separate voting bloc clustering the French-speaking communities in the 2013 and 2014 election cycles, where they are instead grouped in the coastal voting blocs.

3.3. Distribution of question support indicates most voting blocs situated on a gradient with isolated French-speaking communities

Figure 3 shows the structure of support for each referendum question within each voting bloc, with bloc boundaries given by the representative clustering used in Figure 1. For nearly all questions, question support shows a gradient from voting bloc 1 (Old Coast) to voting bloc 5 (Deep Woods),

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

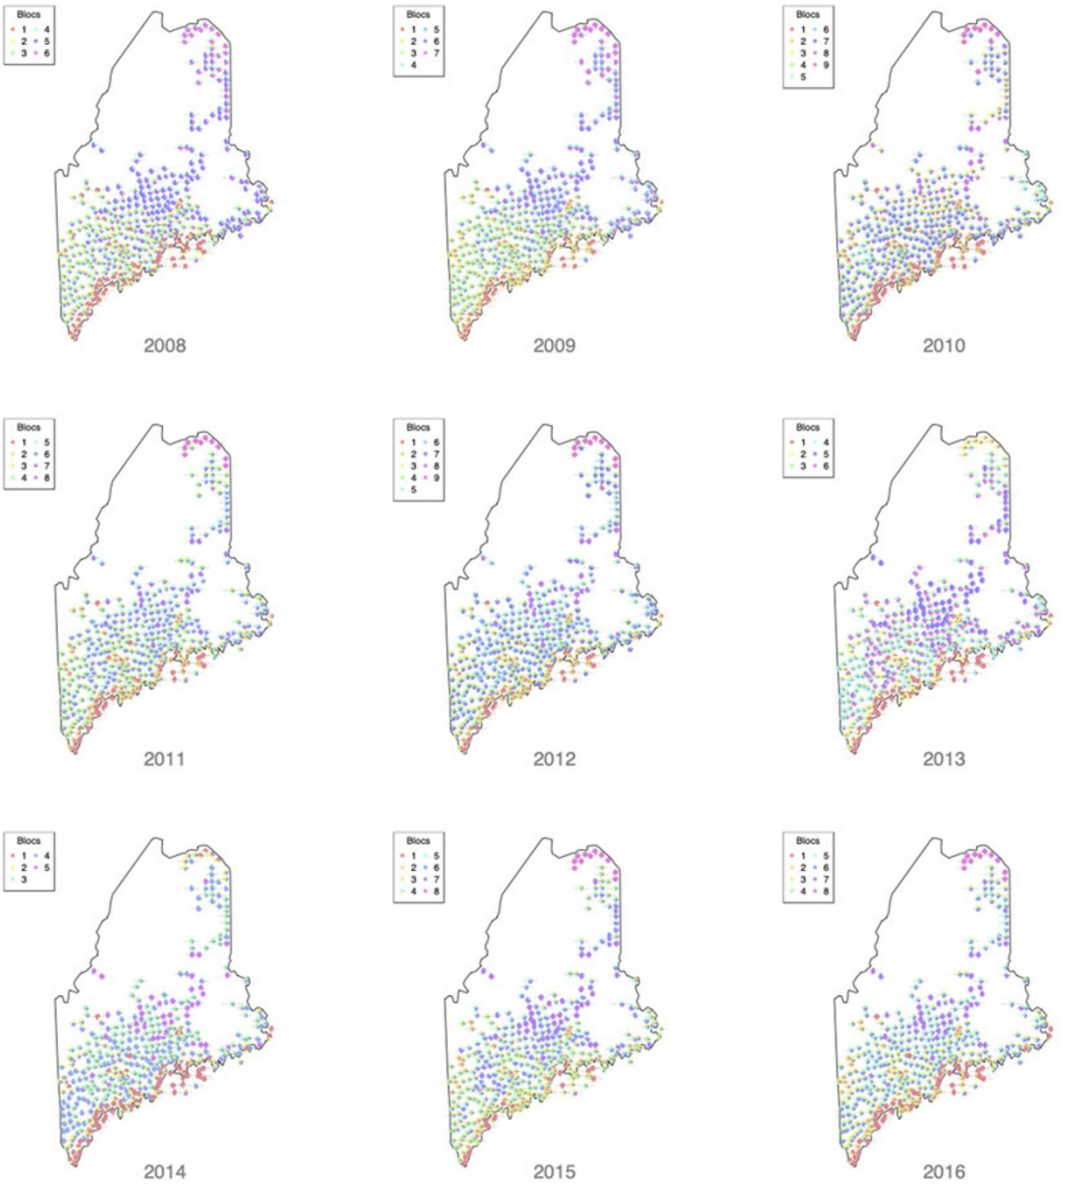


Figure 2. Presentation of representative clusterings for each four-year election cycle from 2008–2011 to 2016–2019. Maps are subscripted with the starting year of the cycle. Number of blocs determined by the posterior mode of K after filtering for clusters smaller than 5 municipalities. Representative clustering determined as in Figure 1(c).

with support often changing incrementally at the bloc boundaries. This leads to a characteristic "bar-code" of question support across the blocs. Five questions (2, 5, 7, 16, 44) show consistent support across the voting blocs. Voting bloc 6, corresponding to the largely French-speaking communities along the state's northern border, is a visible exception to this pattern. This bloc exhibits clustered support for questions out of alignment with the other five blocs, with some questions consistent with voting bloc 5 (largely social/cultural questions such as gay marriage, hunting restrictions, marijuana legalization) while others similar to the Old Coast and New Coast (largely funding-related such as bonds, tax increases for schools, and medical care).

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

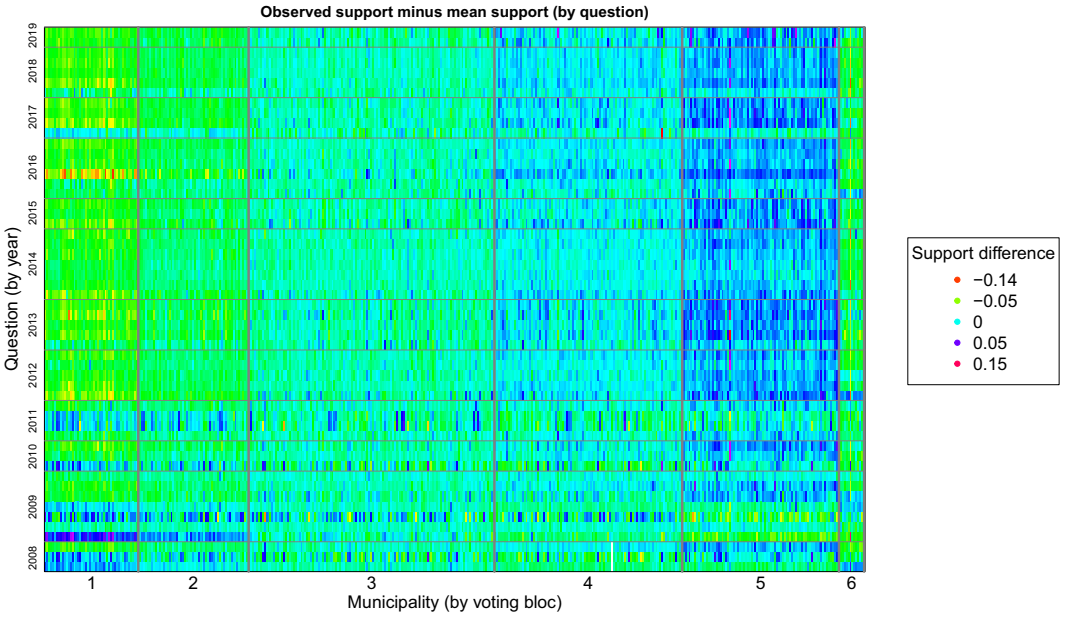


Figure 3. Barcode plot provides a summary of support for each question within each municipality minus the overall level of support for each question across all municipalities. Questions are arranged vertically by year and municipalities arranged horizontally according to the same voting bloc clustering as Figure 1. Grey lines show the boundaries between voting blocs and years.

To understand how these results line up with more traditional methods of assessing clustering in political science, we turn to PCA, frequently used in political science, as well as a *t*-distributed stochastic neighbor embedding (t-SNE), a common technique in machine learning, neuroscience, and genomics Van der Maaten and Hinton (2008); Li et al. (2017) for dimensional reduction. We use a pairwise distance between data entries (here vote totals for each municipality, broken down by question) calculated using a Jensen-Shannon metric added across questions to construct the final distance (Aitchison, 1982). The result is a presentation of each municipality in a 2-dimensional field reflecting the similarity of support across all questions. We then color each of the points according to the mixture proportions determined in Figure 1.

Figure 4 provides two complementary portraits to the barcode plot in Figure 3. In both PCA and t-SNE plots, the first five voting blocs situate in sequence along a (loosely) one-dimensional gradient following party partisanship, with the Old Coast forming the Democratic end, the Deep Woods bloc forming the Republican end, and the remaining blocs filling in the space in between. The t-SNE presentation indicates that there is reduced variation among the data that drive the relatively unmixed communities of the Old Coast and Deep Woods. Notably, the t-SNE separates the unmixed municipalities in the Deep Woods from the mixed communities shared by the Small Mills, even though these communities are often geographically neighboring. Similarly, municipalities identified as mixed between the Hills and Fields and the Small Mills form a boundary in the t-SNE presentation. Most notably, the French-speaking communities in the north of the state form a distinct, separated region in the t-SNE representation that is also indicated by the mixture model. However, this distinction is somewhat ablated in the PCA presentation, which pushes the New France cluster into the New Coast. This makes clear that while PCA and t-SNE can capture much of the broad structure within the data, the mixture model permits identification of more refined local structures.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

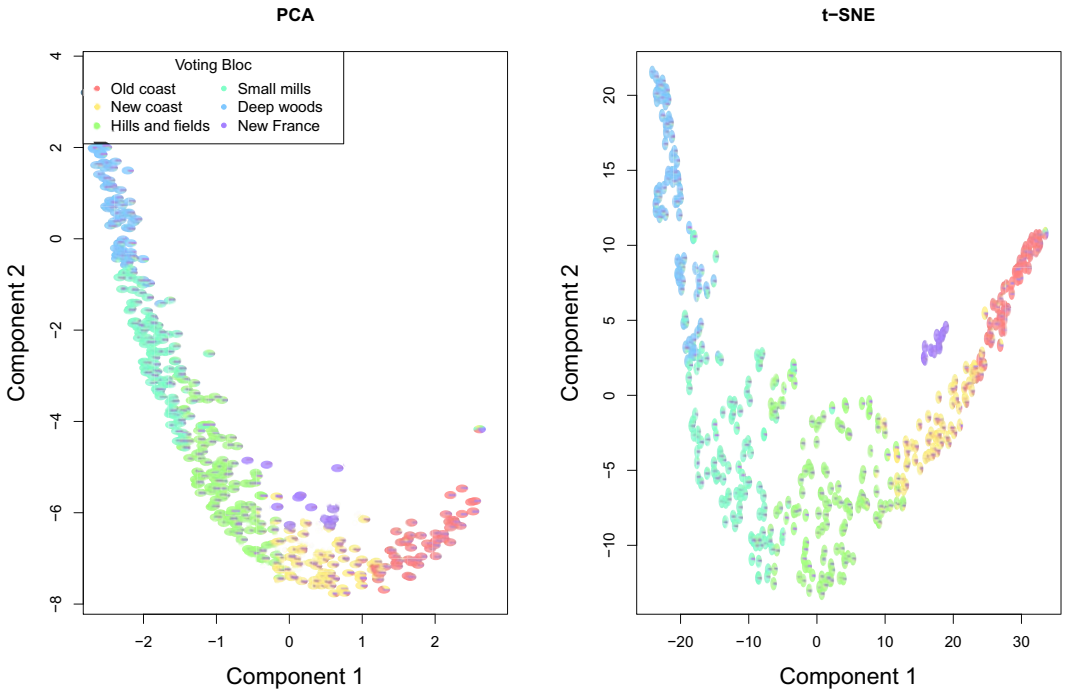


Figure 4. PCA projection and t-SNE embedding in two-dimensions from referendum data using a Jensen-Shannon divergence for pairwise distances. Municipalities are colored according to the mixture proportions used in Figure 1. Colors the same as in Figure 1c.

3.4. *Measures of model fit reveal a small number of "local" questions*

To explore the quality of model fit for individual questions against the complete data set, we compare the posterior predictive distribution against the observed value of support in the data for each question across all municipalities. For each question q and each municipality i, we use the posterior samples of $\alpha_{z_{iq}}$ to build up the posterior predictive distribution of support and compare the resulting distribution's median—$\tilde{p}_{iq}$—against the observed support for the proposition $\hat{p}_{iq}$. For each iteration in the thinned Markov chain, we sample 100 values from Beta distribution according to $(\alpha_{z_i q0}, \alpha_{z_i q1})$. We compare the median of this distribution to the observed proportion ($\hat{p}_{iq} = \frac{c_{iq0}}{c_{iq0}+c_{iq1}}$). Although this statistic neglects variance arising from sample size that would be accounted for with the BB distribution, it makes for an accessible comparison across municipalities of different sizes. We refer to $\tilde{p}_{iq} - \hat{p}_{iq}$ as the ***estimated question fit***.

Figure 5 shows the median and standard deviation of the estimated question fit taken across all municipalities. Nearly all questions exhibit similar values in both statistics, indicating that model fit is consistent across a wide range of municipalities. Notably, five questions show markedly higher standard deviation (in red in Figure 5). Four of the five of these questions concern the creation or regulation of specific casinos within the state. The final question regards the repeal of a school consolidation proposal particularly affecting Maine's rural schools. These "local" questions are natural candidates for violating the nonlocality assumption of the model as voters may have been motivated by specific local concerns that overrode any broader cultural positions. We further examine the performance of each question across all municipalities in Supplementary Figures 1–3 where we plot the estimated question fit for each municipality organized by voting bloc. Within questions, we observe

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

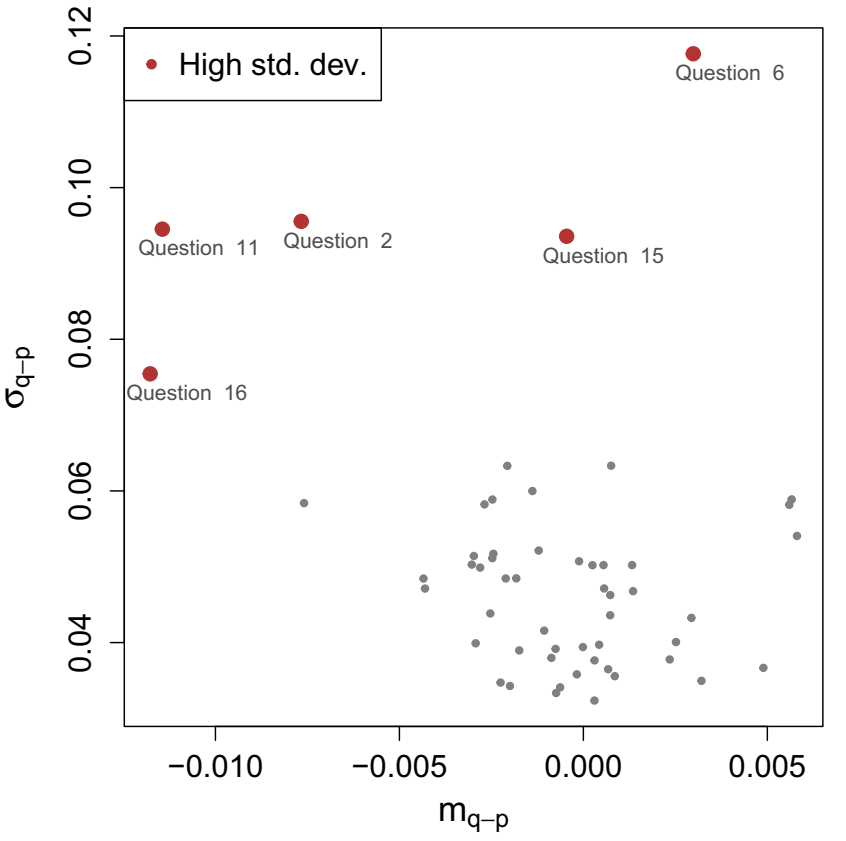


Figure 5. Median and standard deviation of the estimated question fit $(q - p)$ across all municipalities by question. Five questions exhibit higher standard deviation, marked in red. Plots of the estimated question fit by question for each municipality are given in Supplementary Figures 1–3.

strong consistency in the variation across municipalities, with "local" questions exhibiting higher variation across all municipalities and consistent across voting blocs.

To explore the locality of these questions, we plot the estimated fit for each question for each municipality onto a map of Maine (Supplementary Figure 4). For forty-nine questions, the estimated question fit appears to have little geographical structure other than slightly higher values for rural municipalities likely owing to small sample sizes. The five "local" questions exhibit marked geographical structure, as shown in Figure 6. Question 2 posed the creation of a casino in Oxford county, in the western part of the state, where support was inflated relative to the model's expectations. In the far eastern part of the state, support was suppressed relative to the model expectations. This geographical structure justifies the "local" quality of this question: the model assumes that the voting blocs are shared across all municipalities in the state; for questions of concern to specific regions, the model consequently fails to capture that variation, leading to poor fit in a geographically structured fashion.

To further understand the interrelationship between the "local" questions and the voting bloc structure, we plot the Jensen-Shannon distance between each pair of municipalities' observed level of support for each question, organized by the representative clustering (e.g. Supplementary Figure 4) (Briët and Harremoës, 2009). This provides a measure of distance between two probability distributions. We observe that nearly all questions exhibit strong blocking structure where more municipalities with similar support values are gathered according to the representative clustering.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

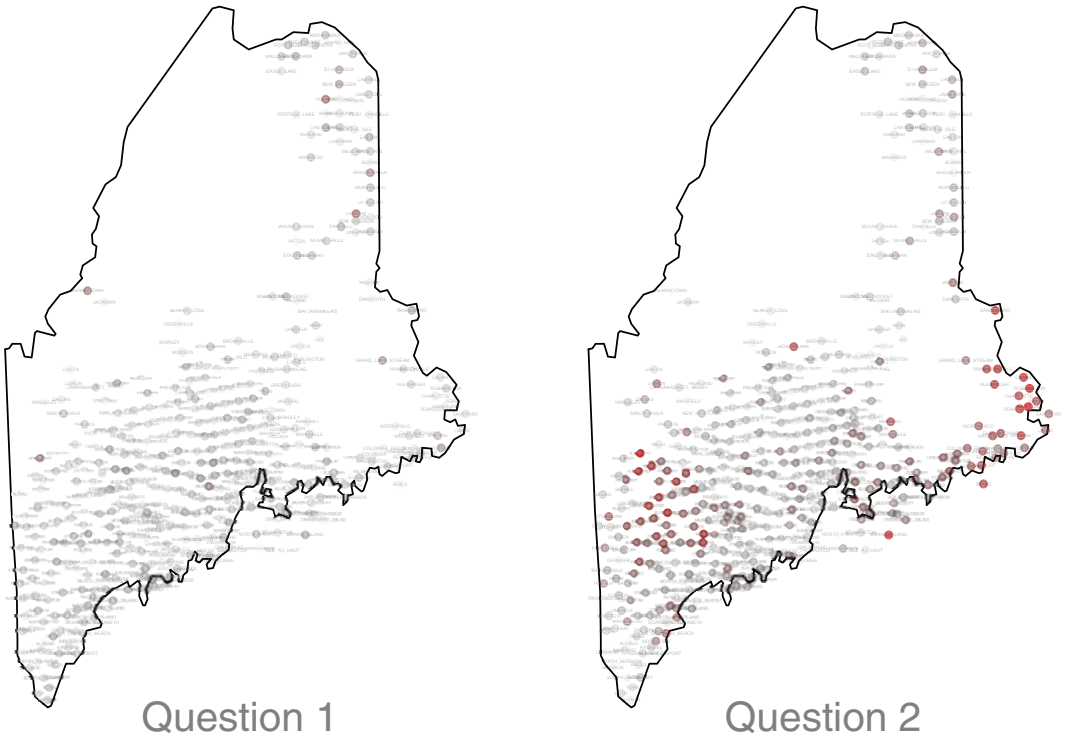


Figure 6. Model fit across a typical question (left) and a "local" question (right). Quality of fit is plotted in a grey to red scale with white being perfect fit (0) and red being the worst observed fit (0.06).

In contrast, the five high variance questions exhibit no such blocking but instead show consistent variability in distance between municipalities within the same bloc. A small number of lower variance questions also exhibit weak blocking: questions 7, 10, 23, and 38. These questions related to governmental functioning. These questions had less separation between the voting blocs in terms of their support, meaning that support was more homogeneous across municipalities, leading to weaker apparent blocking.

4. Discussion

4.1. The model limitations and possible extensions

The primary purpose of this work is to expand the statistical scope of voting bloc models to longitudinal referendum vote totals aggregated at the municipal level. Since individual referendum ballots have only limited information, the model requires multiple ballots to provide stable inference, limiting the applicability to locations where regular referendums occur. However, the frequency of referendums in many polities offers this model ample contexts for application. Many US states make regular use of these elections in a manner comparable to the case study considered here. Switzerland—alone responsible for a third of referendums worldwide—seems particularly attractive as a site of study.

Similar models to the one developed here are widely used in genetics (variants of latent Dirichlet allocation models (Blei et al., 2003; Falush et al., 2003)), microbiome analysis (Dirichlet-multinomial mixture models and phylogenetic mixture models (Holmes et al., 2012; Mao and Ma, 2022)), and text analysis (DiMaggio et al., 2013; Bohr and Dunlap, 2018). In all these contexts, the mixture model serves to capture empirical structures present in the data with limited assumptions

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

about their structure. The broader analytic work is to then explain the observed cluster structure in terms of external variables (for instance, migration, selection, and drift for genomic data; ecological and environmental changes for microbiome data). The model presented here and the approaches of Gormley and Murphy, (and others) lay out a similar framework for political analysis: by identifying structurally similar patterns of voting behavior, theoretical analysis can be calibrated to account for broad trends across longitudinal voting. The analogy is more than just conceptual: the extensive literature on modeling the connections between complex latent processes and the observed mixture structure provides numerous avenues for further research (Hellenthal et al., 2014; Lawson et al., 2018).

A secondary purpose of the paper is to show how this model can be used to uncover spatial and temporal patterns within the data set. In the case study, the election cycle analysis indicates broad consistency in the voting bloc structure across the study period. Similarly, the correlation of mixture proportions across municipalities both in the complete data and the election cycle data shows strong spatial structure. That this spatial structure maps closely to PCA and t-SNE data projections provides further corroboration of the model's inferential capacity. However, neither of these types of variation is directly modeled in this approach. Building out techniques that explicitly account for these covariance structures makes for natural paths for future model development. Again following from Gormley and Murphy (2008b), a mixture-of-expert approach appears to be a straightforward extension to the current model to include covariate information. In the context of temporal variation, since the same group of voters necessarily votes on each ballot, a mixture model that blocks the data by ballot would provide additional flexibility in capturing the temporal structure within the data. Such a model could integrate over the hidden state (the voting bloc) proportional to an underlying mixture structure biased by the probability of that voting bloc showing up to the polls. This would permit direct modeling of the temporal changes in apparent voting bloc proportions within a municipality across time. Such a model extension would also allow for fluctuations in voting bloc turnout to be both estimated and automatically accounted for as part of the inference process.

Another important point for further statistical development is modeling of the correlation among ballot questions. Unsurprisingly, transportation bond questions elicit similar voting patterns. Ballots concerning bonds of all types are generally more similar than cultural questions like gay marriage or marijuana legalization. While ideally the model would take account of the full Q^2 covariance matrix, approaches like the Dirichlet-tree mixture model (Mao and Ma, 2022) may prove more feasible. However, unlike current approaches to using Dirichlet-tree mixture models that assume a fixed phylogeny, such an extension would also require estimation of the underlying tree. In this regard, the problem shows notable similarities to statistical estimation in microbial communities (O'Brien et al., 2014). A distinct but related consideration is the degree that questions themselves adhere to the voting bloc assumptions. While the large majority of questions (49) appear consistent with the model's expectations, for the five questions identified in the Results, we observe spatially clustered deviations (e.g. Supplementary Figure 4). The unusually "local" character of these questions is likely related to the content of the questions: four of the five sought to place casinos in those areas. Adjusting the model to infer and model these questions separately would reduce variance in estimates of voting bloc structure.

In addition to the purposes above, this model introduces to the political science literature the use of a BD-MCMC methodological technique to provide a posterior distribution on the number of voting blocs. Unlike frequently used decision criteria, such as AIC or BIC, that select a single "best" number of components, the BD-MCMC provides a posterior distribution on the number of components that allows for a more coherent representation of the model's uncertainty. This is similar to reversible jump MCMC or thermodynamic integration approaches, but with a reduced computational burden. The relative flexibility of BD-MCMC algorithms also means this technique can likely be re-adapted to explicit models of spatial, temporal, and covariate structure.

In a broader context, referendum data can also be seen as an unusually regular example of the multiview problem, a statistical challenge commonly found in machine learning, medicine, and

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

genomics, that occurs when attempting to reconcile multiple different types of assays (or "views") taken on a single object Yi et al. (2005); Chao et al. (2017); Carmichael (2020). In the context here, the views are created by the different ballot questions that might separate voters along different lines (e.g. transport bonds versus marijuana legalization). This is at least partially reflected in the variable blocking structure observed in Supplementary Figure 4. Unlike many multiview contexts where the views have different distributional structures, referendum data (and likely voting data more broadly) possess a consistent underlying distribution across views. Future model development can leverage this comparison to import frameworks developed for more heterogeneous contexts. In a complementary fashion, referendum data may also serve as a useful training set for multiview algorithm development.

4.2. Connections to other methods in political science

For the political scientist, a natural question is how credibly to interpret the results of the mixture model: how should they understand the components of this model as corresponding to well-defined political dispositions? Do the underlying components represent previously hidden but meaningful subpopulations that have verifiable properties, or are they purely probabilistic components to aid in the empirical characterization of the data's structure? While the tight coupling between the mixture model and the PCA/t-SNE analyses indicates that meaningful latent structure can be extracted from these data, what remains for political science is to determine how the persistence of these structures corresponds to other analytic dimensions. Granting that the mixture model identifies some consistent aspects of a latent space, the broader interpretive issue then relates to three questions: (1) what is the topography of the underlying latent space? (2) how variable is this topography in structure (for instance, in time)?, and (3) how fixed are communities' relation to this topography? At this stage of analysis, there are no certain answers, and indeed they will likely vary by empirical context. In a stationary population like Maine, the latent topography may be relatively fixed both in structure and in the mixture proportions of communities. In locations like California with large internal and external migration, variability of both the structure of the latent space and the composition of communities may rapidly change.

The approach developed here connects to a number of different existing techniques in political science, including ecological regression, polling, and latent class analysis. Ecological regression resolves questions of how well regression is preserved under aggregation, and has been shown to be consistent under specific conditions (King et al., 2004). The approach presented here provides a new way to consider both aggregated vote data and methods for regression (the mixture-of-experts model). An important open statistical question is then how these two approaches can be reconciled, specifically how the consistency conditions of ecological regression can be translated into this new context.

The method presented may also bear upon polling methods. Statistically, much of polling practice focuses on (1) achieving a sufficiently large random sample and (2) unskewing the results to account for demographic categories that are not well-represented in the sample. The presence of latent structure identified by the mixture model may pose a new approach to both the initial procedure of sampling and to the procedures for unskewing. To achieve a representative sample, using the inferred voting blocs' spatial configuration (for instance, unmixed municipalities) to direct the sampling procedure may prove more effective than randomly sampling, since the structure of opinion may be largely accounted for in separation among the voting blocs. Similarly, the inference of voting blocs permits a new type of unskewing procedure: if the respondents' municipalities are recorded, the poll can be unskewed by reweighting the relative presence of the voting blocs in the sample.

4.3. Data extensions

Referendum data from Maine exists from 1910 onward, though it was employed with variable consistency over the 20^{th} century. A natural avenue of research is to develop a dataset that probes further

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

into the past. Maine's election records go back with some consistency to the middle of the 19th century and with modern levels of precision from 1880 onward. From the late 1970s, referendums have been used with increasing frequency and so for at least this period—1970–2022—these data can be used to examine the stability of voting blocs over the period. While the voting blocs are apparently consistent over the time period studied here. Observing consistency (or changes) over a longer time span would give further indications about the use of time blocs in stationary or changing demographic contexts.

Another point for consideration is how similar models might be extended to make inference from other types of election data, most notably candidate elections. Candidate elections in contexts of single-district plurality can likely be accessed directly from the referendum model since a two-candidate contest is similar to a yes/no vote on a referendum. Other forms of candidate elections commonly used in parliamentary systems are more analogous to (although distinct from) RCV data and will likely require careful distributional analysis to develop adequate models. However, the promise is substantial: the strong covariance information present in these ballots could provide information about the structure of the underlying voting blocs.

The extension of models of these types to multiple election types opens up a new frontier in political analysis, where multiple types of elections—referendum, candidate, RCV—may be both simultaneously and complementarily considered. Importantly, these analyses may include different scales of aggregation—municipality, district, state—that can inform understandings of voting blocs and their consistency in time and space. The synthesis of multiple spatial scales and election types across time would present a new framework for contextualizing political opinion both its historical progression and in the present.

Even without the extensions above, the model can be applied to a wide variety of referendum data. In addition to Maine, a number of other US states regularly engage in referendums. For instance, California had roughly three times as many ballot questions as Maine during the study period (although without the regularity of municipal boundaries or stable population as in Maine). Switzerland has a similar number of referendums aggregated at the community level going back to 1980 with digital records, and into the 19th century in analog form.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/psrm.2025.10050. To obtain replication material for this article, https://doi.org/10.7910/DVN/S3DODE.

Funding statement. No funding support.

Competing interests. None.

References

Ahlquist JS and Breunig C (2012) Model-based clustering and typologies in the social sciences. *Political Analysis* **20**, 92–112.

Aitchison J (1982) The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological)* **44**, 139–160.

Amelio A and Pizzuti C (2012) Analyzing voting behavior in Italian parliament: Group cohesion and evolution. In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. IEEE, pp. 140–146.

Amirkhani M, Manouchehri N and Bouguila N (2021) Birth–death MCMC approach for multivariate beta mixture models in medical applications. In Advances and trends in artificial intelligence. Artificial intelligence practices: 34th International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (IEA/AIE 2021), Kuala Lumpur, Malaysia, July 26–29, Proceedings, Part I (pp. 285–296). Springer.

Archives MS (2022) Historical election records. https://www.maine.gov/sos/arc/research/index.html.

Bailey MA, Strezhnev A and Voeten E (2017) Estimating dynamic state preferences from United Nations voting data. *Journal of Conflict Resolution* **61**, 430–456.

Bakk Z, Oberski DL and Vermunt JK (2014) Relating latent class assignments to external variables: Standard errors for correct inference. *Political Analysis* **22**, 520–540.

Beck N and Jackman S (1998) Beyond linearity by default: Generalized additive models. *American Journal of Political Science* 596–627.

Beck NL (2000) Political methodology: A welcoming discipline. *Journal of the American Statistical Association* **95**, 651–654.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

Blackwell M (2018) Game changers: detecting shifts in overdispersed count data. *Political Analysis* **26**, 230–239.
Blei DM, Ng AY and Jordan MI (2003) Latent Dirichlet allocation. *Journal of Machine Learning Research* **3**, 993–1022.
Bohr J and Dunlap RE (2018) Key topics in environmental sociology, 1990–2014: Results from a computational text analysis. *Environmental Sociology* **4**, 181–195.
Bowler S, Donovan T and Tolbert CJ (1998) *Citizens as legislators: Direct Democracy in the United States*. Columbus, OH, USA: The Ohio State University Press.
Briët J and Harremoës P (2009) Properties of classical and quantum Jensen-Shannon divergence. *Physical Review A* **79**, 052311.
Chao G, Sun S and Bi J (2021) A survey on multi-view clustering. *IEEE Transactions on Artificial Intelligence* **2**, 146–168.
Carmichael I (2020) Learning sparsity and block diagonal structure in multi-view mixture models. (ArXiv Preprint). arXiv. https://arxiv.org/abs/2012.15313.
Clinton J, Jackman S and Rivers D (2004) The statistical analysis of roll call data. *American Political Science Review* **98**, 355–370.
Converse PE (1987) Changing conceptions of public opinion in the political process. *The Public Opinion Quarterly* **51**, S12–S24.
D'Angelo S, Murphy TB and Alfò M (2019) Latent space modelling of multidimensional networks with application to the exchange of votes in Eurovision song contest. *The Annals of Applied Statistics* **13**, 900–930.
De Leeuw J (2006) Principal component analysis of senate voting patterns. *Real Data Analysis* 405–411.
DiMaggio P, Nag M and Blei D (2013) Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of US government arts funding. *Poetics* **41**, 570–606.
Dubin JA and Gerber ER (1992) Patterns of voting on ballot propositions: A mixture model of voter types. *Technical report*.
Falush D, Stephens M and Pritchard JK (2003) Inference of population structure using multilocus genotype data: linked loci and correlated allele frequencies. *Genetics* **164**, 1567–1587.
Frühwirth-Schnatter S (2006) *Finite Mixture and Markov Switching Models*, Volume 425. New York, NY, USA: Springer.
Fruhwirth-Schnatter S, Celeux G and Robert CP (2019) *Handbook of Mixture Analysis*, New York, NY, USA: CRC Press.
Gallup G and Rae SF (1940) *The pulse of democracy : the public-opinion poll and how it works.* New York, NY, USA: Simon and Schuster.
Gelman A, Park DK, Ansolabehere S, Price PN and Minnite LC (2001) Models, assumptions and model checking in ecological regressions. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **164**, 101–118.
Gormley CI and Frühwirth-Schnatter S (2019) *Mixture of experts models*. In *Handbook of Mixture Analysis*, Boca Raton, FL: Chapman and Hall/CRC, pp. 271–307.
Gormley CI and Murphy TB (2006) *A latent space model for rank data*. In Edoardo Airoldi, David M. Blei, Stephen E. Fienberg, Anna Goldenberg, Eric P. Xing, Alice X. Zheng, *ICML Workshop on Statistical Network Analysis*, New York, NY: Springer, pp. 90–102.
Gormley CI and Murphy TB (2008a) Exploring voting blocs within the Irish electorate: A mixture modeling approach. *Journal of the American Statistical Association* **103**, 1014–1027.
Gormley CI and Murphy TB (2008b) A mixture of experts model for rank data with applications in election studies. *The Annals of Applied Statistics* **2**, 1452–1477.
Gormley CI and Murphy TB (2010) *Clustering ranked preference data using sociodemographic covariates*. In Stephane Hess and Andrew Daly, *Choice Modelling: The State-of-the-Art and the State-of-Practice*, Emerald Group Publishing Limited.
Green PJ (2003) Trans-Dimensional Markov Chain Monte Carlo. *Oxford Statistical Science Series* 179–198.
Grimmer J, Marble W and Tanigawa-Lau C (2022) Measuring the contribution of voting blocs to election outcomes, 87(3). OSF.
Gunn JA (1995) Public opinion in modern political science. *Political Science in History: Research Programs and Political Traditions* 99–122.
Hellenthal G, Busby GB, Band G, Wilson JF, Capelli C, Falush D and Myers S (2014) A genetic atlas of human admixture history. *Science* **343**, 747–751.
Hill JL and Kriesi H (2001) Classification by opinion-changing behavior: A mixture model approach. *Political Analysis* **9**, 301–324.
Hix S, Noury A and Roland G (2006) Dimensions of politics in the European Parliament. *American Journal of Political Science* **50**, 494–520.
Holloway S (1990) Forty years of United Nations General Assembly voting. *Canadian Journal of Political Science/Revue Canadienne de Science politique* **23**, 279–296.
Holmes I, Harris K and Quince C (2012) Dirichlet-multinomial mixtures: Generative models for microbial metagenomics. *PloS One* **7**, e30126.
Katz JN and King G (1999) A statistical model for multiparty electoral data. *American Political Science Review* **93**, 15–32.
Kim SS (2020) *Three Essays in the Dynamics of Political Behavior*. Pasedena, CA, USA: California Institute of Technology.
Kinder DR (1998) Communication and opinion. *Annual Review of Political Science* **1**, 167–197.
King G (1988) Statistical models for political science event counts: Bias in conventional procedures and evidence for the exponential Poisson regression model. *American Journal of Political Science* **32**, 838–863.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

King G (1990) On political methodology. *Political Analysis* **2**, 1–29.
King G, Tanner MA and Rosen O (2004) *Ecological Inference: New Methodological Strategies*, Columbus, OH, USA: Cambridge University Press.
Koseki SA (2018) The geographic evolution of political cleavages in Switzerland: A network approach to assessing levels and dynamics of polarization between local populations. *PLoS ONE* **13**, e0208227.
Lanza ST, Coffman DL and Xu S (2013) Causal inference in latent class analysis. *Structural Equation Modeling: A Multidisciplinary Journal* **20**, 361–383.
Lawson DJ, Van Dorp L and Falush D (2018) A tutorial on how not to over-interpret STRUCTURE and ADMIXTURE bar plots. *Nature Communications* **9**, 1–11.
Leeper TJ and Slothuus R (2014) Political parties, motivated reasoning, and public opinion formation. *Political Psychology* **35**, 129–156.
Li W, Cerise JE, Yang Y and Han H (2017) Application of t-SNE to human genetic data. *Journal of Bioinformatics and Computational biology* **15**, 1750017.
Magidson J, Vermunt JK and Madura JP (2020) *Latent Class analysis*, Thousand Oaks, CA, USA: SAGE Publications Limited.
Magyar ZB (2022) What makes party systems different? a principal component analysis of 17 advanced democracies 1970–2013. *Political Analysis* **30**, 250–268.
Maine Bureau of Corporations, Elections, and Commissions (2021) *Election results*. Augusta, ME, USA. https://www.maine.gov/sos/cec/elec/.
Mantegazzi D (2021) The geography of political ideologies in Switzerland over time. *Spatial Economic Analysis* **16**, 378–396.
Mao J and Ma L (2022) Dirichlet-tree multinomial mixtures for clustering microbiome compositions. *The Annals of Applied Statistics* **16**, 1476–1499.
May JD (1973) Opinion structure of political parties: the special law of curvilinear disparity. *Political Studies* **21**, 135–151.
McAlister K (2020) *Essays on Latent Variable Models and Roll Call Scaling*. Ph. D. thesis, University of Michigan. HTML link.
McCutcheon AL (1987) *Latent Class analysis*, Thousand Oaks, CA, USA: Sage. p. 64.
McLachlan GJ and Basford KE (1988) *Mixture Models: Inference and Applications to Clustering*, Volume 38. New York: M. Dekker.
McLachlan GJ, Lee SX and Rathnayake SI (2019) Finite mixture models. *Annual Review of Statistics and Its Application* **6**, 355–378.
Mendelsohn M and Parkin A (2001) *Introduction: referendum democracy*. In *Referendum Democracy*, New York, NY: Springer, pp. 1–22.
Mohammadi A, Abegaz F, Heuvel E and Wit E C (2017) Bayesian Modelling of Dupuytren Disease by Using Gaussian Copula Graphical Models. *Journal of the Royal Statistical Society Series C: Applied Statistics* **66**, 629–645. https://doi.org/10.1111/rssc.12171
O'Brien JD, Didelot X, Iqbal Z, Amenga-Etego L, Ahiska B and Falush D (2014) A Bayesian approach to inferring the phylogenetic structure of communities from metagenomic data. *Genetics* **197**, 925–937.
Park JH and Yamauchi S (2023) Change-point detection and regularization in time series cross-sectional data analysis. *Political Analysis* **31**, 257–277.
Pennings P, Kleinnijenhuis J and Keman H (2005) *Doing research in political science: An introduction to comparative methods and statistics*. Thousand Oaks, CA: Sage.
Ratkovic MT and Eng KH (2010) Finding jumps in otherwise smooth curves: Identifying critical events in political processes. *Political Analysis* **18**, 57–77.
Reuning K, Kenwick MR and Fariss CJ (2019) Exploring the dynamics of latent variable models. *Political Analysis* **27**, 503–517.
Schuessler AA (1999) Ecological inference. *Proceedings of the National Academy of Sciences*. **96**, 10578–10581.
Scontras CA (2016) Initiative and referendum: A Maine odyssey. *Lewiston Sun Journal* **10**.
Shi J, Murray-Smith R and Titterington D (2002) Birth-death MCMC methods for mixtures with an unknown number of components. Technical report, Citeseer.
Simon AF and Xenos M (2004) Dimensional reduction of word-frequency data as a substitute for intersubjective content analysis. *Political Analysis* **12**, 63–75.
Spirling A (2007) "Turning points" in the Iraq conflict: Reversible jump Markov chain Monte Carlo in political science. *The American Statistician* **61**, 315–320.
Stephens M (2000) Bayesian analysis of mixture models with an unknown number of components-an alternative to reversible jump methods. *Annals of Statistics* **28**, 40–74.
Van der Maaten L and Hinton G (2008) Visualizing data using t-SNE. *Journal of Machine Learning Research* **9**, 2579–2605.
Vejdemo-Johansson M, Carlsson G, Lum PY, Lehman A, Singh G and Ishkhanov T (2012) The topology of politics: Voting connectivity in the US house of representatives. In *NIPS 2012 Workshop on Algebraic Topology and Machine Learning*.
Wang N, Briollais L and Massam H (2020). The scalable birth–death MCMC algorithm for mixed graphical model learning with application to genomic data integration. https://arxiv.org/abs/2005.04139
Weakliem DL and Biggert R (1999) Region and political opinion in the contemporary united states. *Social Forces* **77**, 863–886.

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press

Weidlich W (1994) Synergetic modelling concepts for sociodynamics with application to collective political opinion formation. *Journal of Mathematical Sociology* **18**, 267–291.

Woodard C (2012) *American Nations: A History of the Eleven Rival Regional Cultures of North America*, Penguin.

Xu S, Ji C and Guedes Soares C (2022) A semiparametric Bayesian method with birth-death Markov Chain Monte Carlo algorithm for extreme mooring tension analysis. *Ocean Engineering* **260**, 111765. https://doi.org/10.1016/j.oceaneng.2022.111765

Yi X, Xu Y and Zhang C (2005) Multi-view EM algorithm for finite mixture models. In Sameer Singh, Maneesha Singh, Chid Apte, Petra Perner, *International Conference on Pattern Recognition and Image Analysis*. Springer, pp. 420–425.

Cite this article: O'Brien J (2026) A Bayesian mixture model captures temporal and spatial structure of voting blocs within longitudinal referendum data. *Political Science Research and Methods*, 1–22. https://doi.org/10.1017/psrm.2025.10050

https://doi.org/10.1017/psrm.2025.10050 Published online by Cambridge University Press