Multidimensional Visualization of Large-Scale Survey Data

Jingtian Yu [1,*], Yanming Di [1], Helen Chang [2], Sharmodeep Bhattacharyya [1], Shirshendu Chatterjee [3] and Neil Hwang [4]

[1] *Department of Statistics, Oregon State University, Corvallis, OR, USA*
[2] *Behavioral and Social Sciences, Hostos Community College - CUNY, New York, NY, USA*
[3] *Department of Mathematics, City University of New York, New York, NY, USA*
[4] *Business and Information Systems, Bronx Community College - CUNY, New York, NY, USA*

Correspondence*:
Jingtian Yu
yujingt@oregonstate.edu

ABSTRACT

Large-scale surveys often contain multiple related response dimensions whose joint structure is difficult to understand using traditional one-dimensional summaries. We propose a visualization framework that combines model-based estimation and two-dimensional heatmaps to reveal multidimensional patterns in survey responses. The framework can be paired with a variety of statistical or machine learning models and is designed to facilitate interpretation of complex survey data through subgroup decomposition, temporal comparison, and exploration of joint response structures. We apply the framework to national survey data from 2014 to 2021 and show that the resulting heatmaps reveal multidimensional structure that is difficult to detect from one-dimensional summaries alone, including clustering, polarization, and differences in issue alignment across policy pairs and population groups. The framework is particularly valuable when associations between policy dimensions are weak and their relationships cannot be readily be inferred from one-dimensional summaries. The framework can be implemented with both conventional regression models and flexible machine-learning approaches, with the latter generally providing closer in-sample fit to the observed policy scores. More generally, the proposed framework provides a scalable and interpretable tool for studying interconnected attitudes in multidimensional survey data.

Keywords: multidimensional visualization, random forests, XGBoost, SVR, heatmaps

1 INTRODUCTION

Recent scholarship underscores that data visualization in political science is not a neutral act of presentation but a consequential practice shaping interpretation, belief, and policy engagement. For instance, Tartaglione and de-Wit (2025) demonstrate experimentally that how variability in out-party political views is visualized critically affects misperception: while interval summaries modestly reduce inflated beliefs about partisan extremism, full distribution displays achieve the strongest correction and diminish support for anti-democratic actions. This finding aligns with broader arguments that visualization embeds social and political choices rather than serving as a passive communication device (Nash et al., 2022; Engebretsen and Kennedy, 2020). Empirical work on citizen-state interfaces, such as Hagen et al. (2019), further illustrates how visual analytics of petition data can render public opinion more legible and actionable for policymakers, even while raising challenges of interpretation and alignment. At a conceptual level, both Engebretsen and Kennedy (2020) and Curini and Franzese (2020) highlight the multiple political roles of visualization—as governance tools, communicative resources, and exploratory devices capable of revealing unexpected patterns or diagnosing data problems. Pedagogical and design-oriented contributions complement this perspective: Zigerell (2023) stresses training students to evaluate and improve visualization practices by emphasizing clarity, cognitive accessibility, and the communication of uncertainty, while Segaran and Hammerbacher (2009) emphasize principles of honesty and transparency in graphical design. Taken together, these studies converge on the view that visualization choices are methodologically consequential and normatively significant, shaping not only how evidence is interpreted but also how political issues and public engagement are structured.

One domain where these visualization approaches are especially relevant is the study of political polarization. Research has documented polarization in U.S. politics from multiple perspectives: rising ideological uniformity and partisan animosity (Pew Research Center, 2014), debates over whether change reflects mass extremism or party sorting (Fiorina and Abrams, 2008; Abramowitz and Saunders, 2008), and the emergence of affective polarization driven by partisan alignment (Mason, 2015). Other work highlights reinforcing mechanisms through socio-cognitive biases and media exposure (Wang et al., 2020), agent-based models of feedback loops (Axelrod et al., 2021), and the interplay of inequality and identity divisions (Stewart et al., 2021). Recent conceptual and methodological work on polarization contributes valuable perspectives (Kleinfeld, 2023; Mehlhaff, 2024; Langford, 2023), but most studies focus on single-policy domains. This underscores the need for visualization methods that capture polarization across multiple, interconnected policy areas, offering more holistic insights into the dynamics of division.

A central goal of this paper is to facilitate the examination of joint patterns of policy preferences rather than analyzing each issue separately. Existing studies in political science often focus on a single policy dimension at a time, such as attitudes toward immigration, gun control, or trade, which may overlook how preferences align across multiple issue areas. Policy preferences are inherently multidimensional, and meaningful patterns often emerge through the alignment of attitudes across multiple issue areas rather than within any single policy dimension alone. However, multidimensional response structures are often more difficult to summarize, visualize, and interpret than individual policy dimensions, which has contributed to the continued reliance on one-dimensional analyses. Understanding these joint structures can reveal clustering, heterogeneity, polarization, and issue-alignment patterns that may be difficult to discern from separate one-dimensional summaries.

To address this challenge, we propose a model-based visualization framework that combines model-based estimation with visualization of joint response distributions. The framework is designed to facilitate exploration of multidimensional survey structure across policy domains, demographic groups, and time periods. By providing interpretable visual summaries of estimated response patterns, the approach can be applied systematically across policy pairs, population groups, survey years, and alternative modeling strategies. Our emphasis is on visualizing model-based estimates rather than raw survey responses. To clarify the contribution of the estimation step, the framework contrasts model-based heatmaps directly against raw-score heatmaps. This comparison demonstrates that fitted-score visualizations can uncover joint preference structures that raw frequency plots obscure due to sparsity, sampling variability, or subgroup imbalance. Direct visualization of raw survey data is fundamentally limited to counting frequencies on a static grid, leaving it vulnerable to sampling variability, sparse regions of the response space, and demographic imbalances, particularly when examining joint dimensions. Crucially, it does not incorporate the additional rich, multidimensional demographic covariates that actually shape those attitudes. Model-based estimation, in contrast, incorporates these very covariates into the estimation step and stabilizes patterns by borrowing information across respondents, yielding a covariate-aware representation that offers deeper insight into the joint policy space.

The proposed framework is not tied to any specific estimation method. Instead, it is designed to improve understanding of complex survey data by providing a common visualization framework for examining estimated response structures. The goal is to help researchers identify substantively meaningful patterns that may be difficult to detect from raw data or separate one-dimensional summaries alone.

We illustrate the framework using policy preference data from the Cooperative Election Study (CES), a large-scale national survey that measures attitudes across multiple policy domains and demographic groups. These examples highlight several ways in which the proposed approach can be used to explore multidimensional survey structure. First, the partisan decomposition of healthcare and abortion preferences shows that the aggregate distribution contains a multipolar structure that is not well summarized by a single left–right average or by either policy dimension alone (Section 3.4). Second, the temporal analysis of abortion–trade and immigration–trade preferences shows that both the location of policy clusters and the strength of issue alignment change over time (Section 3.5). Third, the within-party subgroup analysis demonstrates how demographic decomposition can reveal sources of internal heterogeneity, with gender accounting for much of the Republican subgroup structure in the environment–gun control policy space (Section 3.6). Together, these examples illustrate how the framework can be used to identify clustering, issue alignment, temporal change, and subgroup heterogeneity in multidimensional survey data.

The remainder of the paper is organized as follows. We first introduce the Cooperative Election Study (CES) dataset and the key variables used in the analysis in Section 2.1. We then present our multidimensional heatmap framework in Section 2.2, to reveal patterns of ideological alignment across pairs of policy domains. Section 3 reports the empirical findings using the CES dataset from 2014 to 2021. Finally, Section 4 summarizes the methodological contributions of the framework and concludes with notes on reproducibility and directions for future research.

2 MATERIALS AND METHODS

2.1 Datasets for Analysis

We demonstrate our visualization framework by examining trends in political polarization using the Cooperative Election Study (CES) dataset. Formerly known as the Cooperative Congressional Election Study (CCES), CES is a widely used dataset in political and social science research. Utilizing CES data, Kuriwaki et al. (2024) analyze racial voting patterns across U.S. congressional districts; McCabe et al. (2024) show that healthcare-based voter registration initiatives are particularly effective at reaching younger and more racially diverse populations; and Brown et al. (2025) investigate the rural-urban political divide in the U.S. Originally launched in 2006, the CES has become an essential resource for examining political attitudes, voting behavior, and demographic trends in the United States. Each survey cycle includes over 50,000 respondents, providing a highly representative sample suitable for detailed subgroup analyses.

The CES dataset's large sample size and longitudinal coverage from 2006 to 2021 make it well suited for our visualization framework. In this paper, a heatmap refers to a two-dimensional binned display in which the horizontal and vertical axes represent two policy preference scores, and the color intensity of each cell represents the proportion of respondents whose scores fall in that region of the joint policy space. In raw-score heatmaps, these scores are observed policy scores; in model-based heatmaps, they are fitted policy scores obtained from the estimation step. This structure allows us to conduct detailed subgroup analyses and track how polarization patterns change across multiple policy domains over time.

2.1.1 Response variable: policy preference scores

We define policy preference scores from CES survey responses to summarize respondents' positions across several policy domains. The analysis uses the Cumulative CES Policy Preferences dataset (Dagonel, 2021), which combines policy preference questions from CES surveys over time. In this study, we focus on the policy domains listed in Table 1. Each item is coded numerically so that the response can be incorporated into a domain-specific score.

For each policy domain, we construct a raw score by summing the positively oriented components and subtracting the negatively oriented components, as summarized in Table 2. For example, the raw abortion score is constructed as

$$s_{\text{abortion}} = -\text{abortion}_{\text{always}} + \text{abortion}_{\text{conditional}} + \text{abortion}_{\text{20 weeks}} + \text{abortion}_{\text{prohibition}}.$$

The component names in Table 2 refer to the subscripts of the survey-question variables listed in Table 1. This construction aligns the direction of the policy scores so that higher values correspond to more restrictive, interventionist, or regulation-supportive positions within the corresponding policy domain.

Each raw policy score is normalized to the interval $[0, 1]$:

$$S_{\text{policy}} = \frac{s_{\text{policy}} - \min(s_{\text{policy}})}{\max(s_{\text{policy}}) - \min(s_{\text{policy}})}.$$

This normalization makes the policy scores comparable across domains with different numbers of component questions.

Policy domains were included in model fitting only for survey years in which the corresponding component items were available. Policy-year combinations with unavailable items were excluded from the year-specific modeling datasets; these exclusions are listed in Table S1. For available policy domains, respondent-level scores were treated as missing when any required component item for that domain was missing, and no item-level imputation was used. The model-fitting datasets were then restricted to respondents with nonmissing policy scores, demographic covariates, and survey weights.

The additive construction treats the component items within each policy domain as contributing equally after their response directions are aligned. This choice is consistent with prior public-opinion research that uses multiple survey items to construct more reliable measures of policy attitudes and to summarize related issue positions within broader domains (Ansolabehere et al., 2008; Ansolabehere and Jones, 2010). Similar additive or domain-level indices have also been used in empirical studies of political attitudes and policy preferences (Hetherington et al., 2016; Reid et al., 2025; Lindh and Andersson, 2025), including work on gun-policy attitudes using multiple reform items (Hansen and Dolan, 2025). In the present study, the score is designed as a transparent and reproducible domain-level index for visualization, rather than as a

Abortion	
$\text{abortion}_{\text{always}}$	Always allow abortion
$\text{abortion}_{\text{conditional}}$	Allow abortion only under certain cases
$\text{abortion}_{\text{20 weeks}}$	Ban abortion after 20 weeks
$\text{abortion}_{\text{prohibition}}$	Total prohibition of abortion
Environment	
$\text{enviro}_{\text{carbon}}$	Allow EPA to regulate carbon dioxide emissions
$\text{enviro}_{\text{renewable}}$	Require states use a minimum amount of renewable fuels
$\text{enviro}_{\text{airwateracts}}$	Strengthen EPA enforcement of Clean Air and Water acts
Gun Control	
$\text{guns}_{\text{assaultban}}$	Ban assault rifles
$\text{guns}_{\text{permits}}$	Ease ability to obtain concealed-carry permits
Immigration	
$\text{immig}_{\text{legalize}}$	Grant conditional legal status to undocumented
$\text{immig}_{\text{border}}$	Increase border security between US-Mexico
Health Care	
$\text{healthcare}_{\text{aca}}$	Repeal the Affordable Care Act
$\text{healthcare}_{\text{acamandate}}$	Restore ACA mandate requiring everyone to be insured
$\text{healthcare}_{\text{medicare}}$	Provide Medicare for all Americans
Military	
$\text{military}_{\text{oil}}$	Approve of military ensuring oil supply
$\text{military}_{\text{terroristcamp}}$	Approve of military to destroy terrorist camp
$\text{military}_{\text{genocide}}$	Approve of military to intervene in genocide or civil war
$\text{military}_{\text{democracy}}$	Approve of military to assist in spread democracy
$\text{military}_{\text{protectallies}}$	Approve of military to protect US allies under attack
$\text{military}_{\text{helpun}}$	Approve of military to help the United Nations uphold international law
Spending	(1 = Increase, 3 = Maintain, 5 = Decrease)
$\text{spending}_{\text{welfare}}$	Spending preferences on welfare
$\text{spending}_{\text{healthcare}}$	Spending preferences on health care
$\text{spending}_{\text{education}}$	Spending preferences on education
$\text{spending}_{\text{police}}$	Spending preferences on law enforcement
$\text{spending}_{\text{infrastructure}}$	Spending preferences on transportation/infrastructure
Trade	
$\text{trade}_{\text{china}}$	Tariffs on goods imported from China
$\text{trade}_{\text{canmex include}}$	Tariffs on steel and aluminum INCLUDING from Canada and Mexico

Note: Survey question wording may differ across years; see Dagonel (2021) for details.

Table 1. Policy preference questions from CES and brief description

latent-variable measurement model with item-specific loadings or weights. The resulting scores therefore summarize respondents' expressed positions across the available items in each domain using a common scale.

Policy score	Positive components	Negative components
Abortion	conditional, 20 weeks, prohibition	always
Environment	carbon, renewable, airwateracts	–
Gun control	assaultban	permits
Immigration	border	legalize
Health care	acamandate, medicare	aca
Military	oil, terroristcamp, genocide, democracy, protectallies, helpun	–
Spending	welfare, healthcare, education, infrastructure	police
Trade	china, canmex include	–

Note: Positive components enter the corresponding policy score with a positive sign, while negative components enter with a negative sign. Each raw score is then normalized to the interval $[0, 1]$.

Table 2. Construction of policy preference scores

After defining the policy preference scores, we first examine their one-dimensional distributions over time. Figure 1 shows the yearly marginal distributions of policy preference scores from 2014 to 2021 across the policy domains. These plots are useful for checking whether a single policy domain shifts over time and for identifying years with more concentrated or more dispersed marginal distributions.

However, Figure 1 also illustrates the limitations of one-dimensional summaries for the purpose of studying multidimensional polarization. Each panel displays only one policy domain in one year, so the figure separates the policy dimensions rather than showing how they are jointly organized. As a result, it cannot answer questions such as whether respondents with high immigration scores also tend to have high trade scores, whether two policy domains form distinct clusters, or whether polarization emerges through combinations of issue positions rather than through marginal shifts alone. Two populations can have similar one-dimensional distributions but very different joint structures across policy pairs. For this reason, Figure 1 is informative as a descriptive baseline, but it is not sufficient for our main goal of analyzing alignment, clustering, and heterogeneity across multiple policy dimensions.

This limitation motivates the use of two-dimensional heatmaps in the following analysis. By placing two policy scores on the horizontal and vertical axes and using color intensity to summarize concentration in the joint policy space, the heatmaps reveal cross-policy structure that is not visible from the marginal distributions alone.

2.1.2 Explanatory variables

To examine the drivers of polarization using the model-based visualization framework, we incorporate a set of demographic characteristics from the Cumulative CES Common Content dataset (Kuriwaki, 2023), including gender, race, education level, geographic location, and self-identified partisanship. These variables serve as explanatory inputs to our models and are used to generate fitted policy scores for the visualization framework. They provide the basis for examining how multidimensional policy preferences vary across demographic groups and for decomposing observed patterns by subgroup.

Among these, education level and geographic location are treated as ordinal variables. Higher values of education correspond to higher educational attainment levels (e.g., from high school to graduate degree), while higher values of the geographic variable indicate more rural settings. To derive geographic information, we use county-level Federal Information Processing Standards (FIPS) codes to merge the CES data with the 2023 Rural-Urban Continuum Codes (RUCC) from U.S. Department of Agriculture, Economic Research Service (2023). The RUCC classifies counties on a scale that captures both population size and proximity to urban centers, allowing us to systematically assess the rural-urban context of each respondent. An important aspect of longitudinal survey analysis is the consistency of respondent characteristics over time. In this study, we observe that the demographic composition of the survey sample remained relatively stable across multiple years. This stability is advantageous, as it reduces the likelihood that observed changes in policy preferences or attitudes are driven by shifts in the underlying population structure. Instead, it strengthens the interpretation that temporal trends reflect genuine changes in public opinion rather than sampling variability. Consistent demographics also enhance the comparability of model estimates and subgroup analyses across years, allowing for more reliable inference about evolving policy preferences within and between demographic groups.

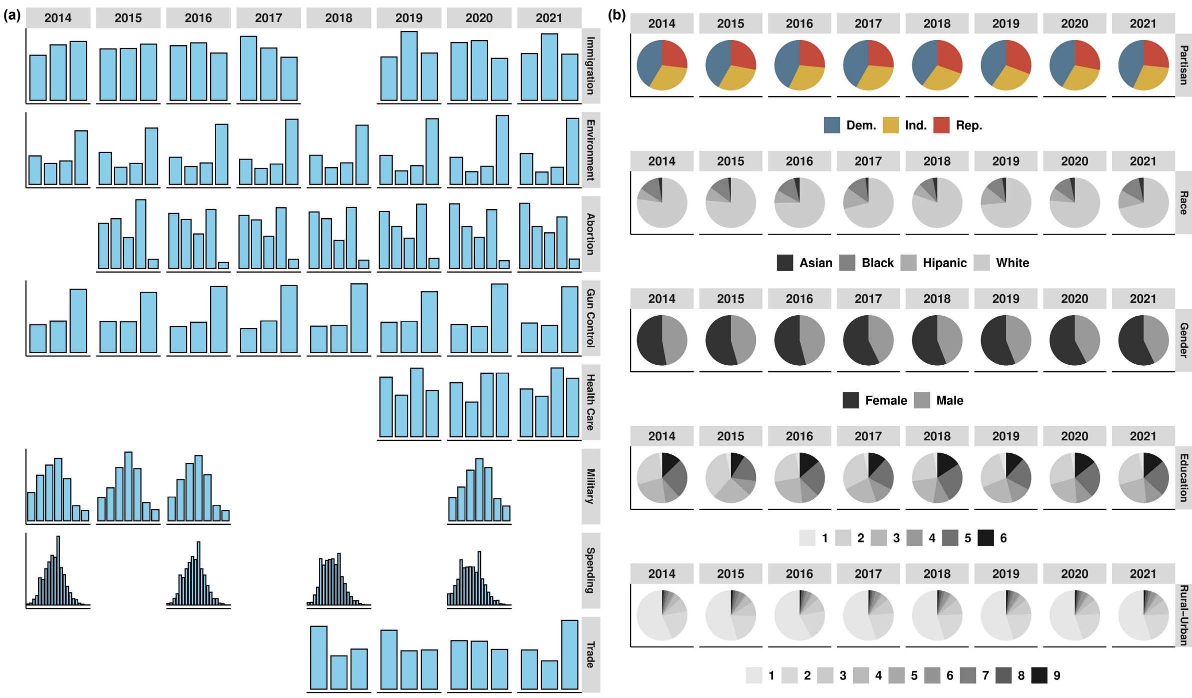

Figure 1. Policy preference distributions and respondent composition. (a) Distribution of policy preference scores by year and topic, 2014–2021. The panels show yearly distributions across four policy domains: abortion, environment, gun control, and immigration. Each row corresponds to a policy topic, and each column corresponds to a survey year. Bar heights represent the proportion of respondents at each score level, allowing comparison of opinion shifts over time. (b) Distribution of explanatory variables from 2014 to 2021. The panels show the yearly composition of five demographic variables used in the analysis: partisanship, race, gender, education level, and rural–urban classification. Each pie chart represents the proportion of respondents in each category for a given year. The distributions are relatively stable across years, indicating limited demographic drift in the dataset and supporting year-to-year comparisons.

2.2 Methods

The goal of the proposed framework is to visualize multidimensional policy preferences in a way that reveals joint structure across policy domains. Direct visualization of raw survey responses is often limited by discrete response scales and sparse combinations of response patterns. To address this issue, we first use demographic characteristics to obtain model-based estimates of policy preference scores and then visualize the joint distribution of these scores using two-dimensional heatmaps. The resulting model-based heatmaps provide a smoother, demographically informed representation of the policy space, facilitating the identification of clusters, issue alignment, subgroup differences, and temporal trends.

2.2.1 Estimation step for policy preference scores

Our visualization framework is built on estimated policy preference scores. For each survey year and policy domain, we use respondents' demographic characteristics to obtain a model-based estimated score. These estimated scores are then used as the input to the two-dimensional heatmaps described below.

The framework is intentionally model-agnostic. Any supervised method that maps respondent characteristics to policy preference scores can be used. In the empirical analysis, we implement four common estimation models: linear models (LM), random forests (RF), support vector regression (SVR), and extreme gradient boosting (XGBoost, abbreviated XGB). These models provide different ways to summarize the relationship between demographic characteristics and policy preferences, but they are not themselves the methodological focus of the paper. Instead, they serve as alternative engines for producing estimated scores that can be visualized within the same graphical framework.

The linear model provides a familiar regression-based benchmark for social science applications. In this setting, each policy preference score is modeled as an additive function of demographic covariates, similar to standard survey analyses that estimate how attitudes vary by partisanship, race, gender, education, geography, and other respondent characteristics. Its main advantage is transparency: the fitted values are generated from a simple and interpretable structure, making the linear model a useful baseline against which more flexible estimation methods can be compared. Random forest is a tree-based ensemble method that can accommodate mixed covariate types and capture nonlinearities and interactions without requiring a prespecified functional form. Support vector regression provides a smoother nonlinear estimation approach through kernel-based regression. XGB is a boosting-based tree ensemble that sequentially improves performance while incorporating regularization. For all four models, the cumulative survey weights are used as observation weights to align the estimation step with the survey design.

The purpose of using estimated scores is to construct a smoothed, demographically informed representation of policy preferences. Furthermore, demographic covariates are incorporated to estimate policy preferences conditional on observed respondent characteristics. The resulting heatmaps therefore summarize systematic demographic structure in the policy space rather than merely smoothing observed response frequencies. Compared with heatmaps based directly on raw response scores, model-based heatmaps reduce sparsity in the two-dimensional policy space and capture systematic variation associated with observed respondent characteristics. This allows the visualization to highlight joint policy structures, temporal changes, partisan decomposition, and subgroup heterogeneity within a common graphical framework.

Because our contribution is the visualization of estimated policy-preference distributions rather than the development or comparison of estimation algorithms, formal model comparison is outside the scope of this study. We use four estimation models only to show that the proposed heatmap workflow can be implemented with both conventional regression models and more flexible machine-learning approaches.

2.2.2 Workflow for constructing heatmaps from model-based estimations

Our visualization framework combines supervised estimation with two-dimensional heatmaps to reveal how policy attitudes co-vary across issues. The workflow proceeds in five steps.

1. **Construct policy-domain responses.** For each survey year, we construct policy preference scores for all available policy domains using the same coding rules applied throughout the data pipeline. If a policy response is unavailable in a given year, that policy-year combination is excluded from model fitting.

2. **Process demographic covariates.** For the same respondents, we assemble a common set of demographic explanatory variables, including variables such as gender, race, education, partisanship, and geographic location. Covariates are processed consistently across all four estimation models so that differences in the final heatmaps reflect differences in model structure rather than differences in input construction.
3. **Fit year-specific estimation models.** For each survey year t and each policy domain d, we fit one of the four models described above to estimate policy scores from demographics:

$$\hat{S}^{(t)}_{id,m} = \hat{f}^{(t)}_{d,m}\left(X^{(t)}_i\right), \qquad m \in \{\mathrm{LM}, \mathrm{RF}, \mathrm{SVR}, \mathrm{XGB}\}.$$

This yields a estimated score for every respondent, policy domain, year, and model. The model-specific estimations preserve a common output structure, which facilitates downstream comparison across modeling approaches.
4. **Construct heatmaps from estimated policy pairs.** For any two policy domains, say d_1 and d_2, we pair the estimated scores

$$\left(\hat{S}^{(t)}_{id_1,m}, \hat{S}^{(t)}_{id_2,m}\right)$$

for all respondents and divide the two-dimensional score space into a regular grid. Each respondent is then assigned to a grid cell according to the pair of estimated scores, and the cell frequency is represented by a color gradient. The resulting heatmap summarizes the joint distribution of model-based policy preferences for that year and model.
5. **Compare heatmaps across years and groups.** Because the estimation step is repeated separately by year and policy domain, the same procedure can be used to compare heatmaps over time and across demographic subgroups. This makes it possible to examine whether clustering, policy bundling, or polarization patterns shift across survey years or vary across population groups. Although the framework can be implemented with different estimation models, our focus is on using these models as alternative ways to generate policy-score estimations rather than on formally comparing model performance.

The key advantage of this workflow is that it provides a common framework for examining the joint distribution of policy preferences across policy domains, demographic groups, survey years, and estimation models. By placing LM, RF, SVR, and XGB within the same visualization pipeline, the framework provides a scalable and flexible tool for studying polarization and multidimensional issue alignment in survey data.

2.2.3 Advantages model-based heatmaps

Traditional raw-score summaries describe the distribution of each policy score separately. These one-dimensional plots are useful for showing marginal shifts over time, but they cannot reveal whether respondents with similar positions on one issue also tend to share positions on another issue. As a result, they may overlook clustering, dispersion, and issue bundling across policy domains.

Model-based heatmaps address this limitation by modeling the full dataset and then visualizing the joint distribution of estimated policy scores. This approach has two main advantages over one-dimensional raw-score distributions. First, it represents two policy dimensions simultaneously, making it possible to detect joint structures such as clusters, diagonal alignment, multipolar patterns, or diffuse distributions that are not visible from marginal plots alone. Second, it incorporates multiple demographic covariates into the estimation step, so the resulting heatmaps summarize systematic variation in policy preferences associated with respondent characteristics such as partisanship, gender, race, education, and geography.

Compared with the raw-score two-dimensional distribution, the model-based heatmap provides a clearer view of the underlying joint organization of policy preferences. The raw two-dimensional plot is useful as a transparent descriptive benchmark, but because the observed scores are discrete, the distribution can appear as a small number of isolated response combinations. Furthermore, the raw display is strictly limited to the two policy dimensions being plotted, ignoring other highly informative individual characteristics that explain the data. The model-based heatmap instead places respondents in a smoother, continuous policy space after incorporating demographic covariates, making it easier to detect diagonal alignment, clustering, multipolar structure, and intermediate groups. A direct comparison is provided in Section 3.2, where the model-based heatmap reveals a three-pole structure that is less apparent in the raw-score display.

In this sense, model-based heatmaps complement the raw-score summaries by providing a more comprehensive view of policy alignment. Rather than focusing only on marginal variation within separate policy domains or on discrete raw-score combinations, they summarize conditional relationships in a common two-dimensional framework, making multidimensional polarization patterns easier to detect and interpret.

3 RESULTS

This section illustrates the proposed model-based visualization framework through a series of examples that progress from simple to more complex applications. We begin by introducing the model-based heatmap and demonstrating how joint policy preferences can be visualized in a two-dimensional policy space. We then compare model-based and raw-score visualizations, examine examples with different levels of policy alignment, and explore how the framework can be used to investigate partisanship, temporal dynamics, and subgroup heterogeneity. Together, these examples illustrate how model-based heatmaps can reveal multidimensional patterns in survey data that are difficult to discern from one-dimensional summaries alone.

3.1 Introducing the model-based heatmap

Figure 2 presents the model-based heatmap for the full population in 2015, based on a RF model, with estimated gun control scores on the horizontal axis and estimated immigration scores on the vertical axis. Each tile represents the proportion of respondents whose estimated scores fall in that bin. The color scale ranges from light blue for very low proportions to yellow, orange, and red for increasingly dense regions, so warmer colors indicate where respondents are most concentrated.

The main feature of the figure is a two-cluster structure. One cluster appears in the upper-middle region, corresponding to relatively high immigration scores and moderate gun-control scores. A second cluster appears in the lower-right region, corresponding to high gun-control scores and lower immigration scores. These two high-density regions suggest that the joint distribution is not centered around a single average profile, but instead organized around two distinct modes.

The marginal histograms and density curves are consistent with this interpretation. Along the bottom, the distribution of estimated gun-control scores is concentrated in the middle-to-high range. Along the left, the immigration distribution shows multiple local concentrations rather than a single peak. Together, the heatmap and marginals show that the full population exhibits a structured, bimodal pattern in joint policy preferences.

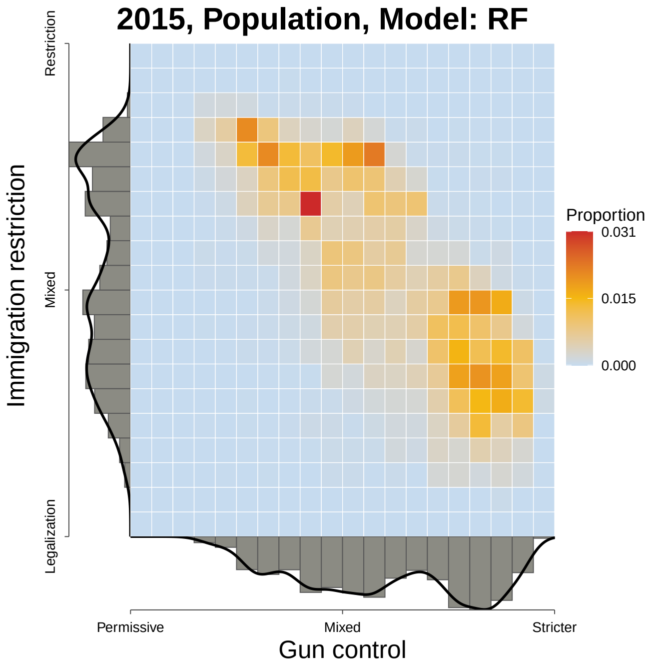

Figure 2. Joint policy preferences in a single year. RF-based heatmap of estimated gun control and immigration scores for the full population in 2015. The horizontal axis shows estimated gun control scores, and the vertical axis shows estimated immigration scores. Cell color represents the proportion of respondents in each two-dimensional bin, with warmer colors indicating higher concentration. The marginal histograms and density curves summarize the corresponding one-dimensional distributions along each policy dimension.

3.2 Comparing raw-score and model-based heatmaps

Figure 3 compares the raw-score visualization and the model-based heatmap for public healthcare and immigration preferences in 2019. In both panels, the horizontal axis represents public healthcare preferences, ranging from market-based to government-led, and the vertical axis represents immigration preferences, ranging from legalization to restriction. The raw-score plot in panel (a) provides a direct descriptive summary of the observed response combinations. Because the raw scores take only a limited number of possible values, the distribution is represented by isolated bubbles on a discrete grid. This display serves as a transparent baseline, but the joint structure is difficult to interpret beyond the most common response combinations. The model-based heatmap in panel (b) shows a more organized joint distribution. Rather than exhibiting a single central cluster or a simple two-pole division, the RF-based heatmap displays three visible concentrations. One concentration appears near the market-based healthcare and restrictive-immigration region, a second appears near the mixed region, and a third appears near the government-led healthcare and legalization region. Together, these concentrations form a downward diagonal pattern, suggesting that healthcare and immigration preferences are jointly structured but not well summarized by a single marginal distribution or a simple binary division.

This comparison highlights the added value of the model-based visualization. The raw-score plot records the observed response frequencies, but the discreteness of the raw scale can obscure broader patterns in the joint policy space. By incorporating respondent characteristics through the estimation step and visualizing the resulting estimated scores in a common two-dimensional space, the heatmap makes the multipolar structure more apparent. In this example, the most notable pattern is the emergence of a

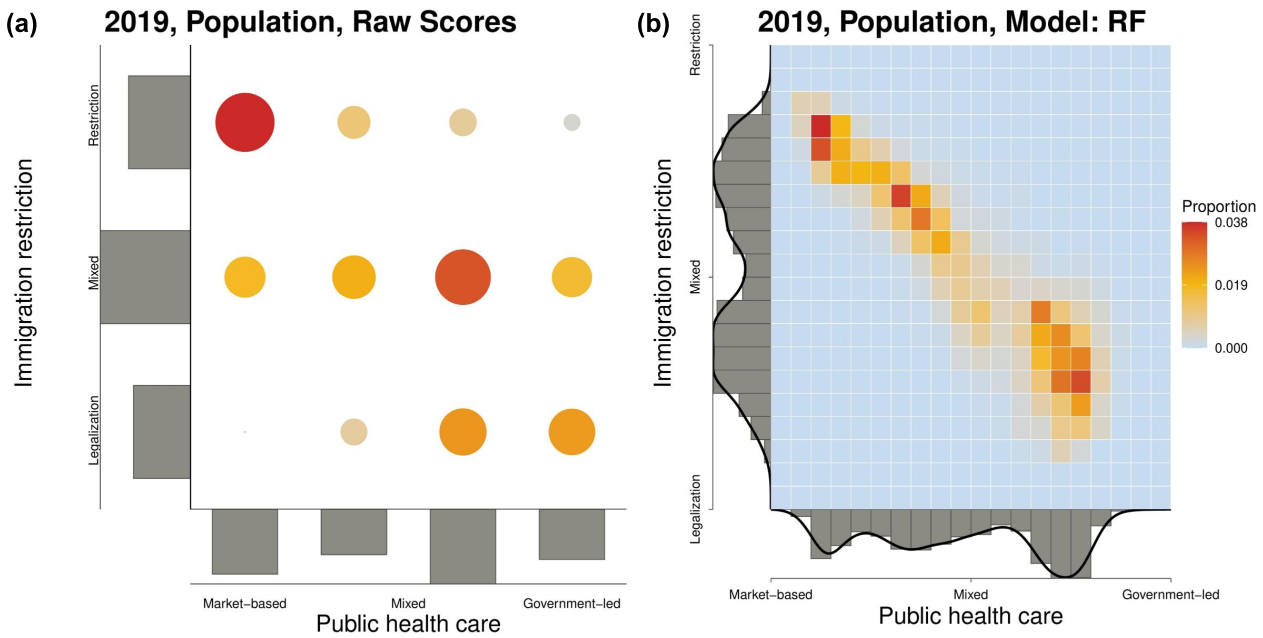

Figure 3. Comparison between raw-score and model-based heatmaps. Joint distribution of public healthcare preferences and immigration restriction preferences for the full population in 2019. Panel (a) shows the raw-score visualization, where the policy scores are observed on a discrete response scale and the dot size and color represent the proportion of respondents in each raw-score combination. Panel (b) shows the RF-based estimation heatmap, where respondents are placed in a continuous two-dimensional policy space using estimated scores. Warmer colors indicate higher concentrations of respondents. The model-based heatmap reveals a clearer three-pole structure that is less apparent in the raw-score display.

three-pole distribution, including an intermediate mixed pole that is less clearly separated in the raw-score display.

The comparison in Figure 3 helps readers see how the proposed framework relates to the observed policy scores while making multidimensional patterns more interpretable than raw-score heatmaps alone. The raw-score heatmaps show the observed distribution of constructed policy scores, while the model-based heatmaps provide smoothed summaries of fitted policy-preference distributions conditional on the covariates included in the model. The main contribution of the framework remains the visualization of policy-preference patterns rather than formal model validation or model selection.

3.3 Comparing high- and low-correlation heatmaps

Figure 4 contrasts two examples that highlight how the heatmap depends on the relationship between the two policy dimensions. In panel (a), the 2018 SVR-based heatmap for abortion and spending shows a strong positive association. Most of the mass is concentrated along a narrow diagonal band, indicating that respondents with higher estimated abortion scores also tend to have higher estimated spending scores. In this case, the heatmap mainly reinforces a pattern that is already suggested by the marginal distributions.

Panel (b) tells a different story. The 2021 XGB-based heatmap for immigration and trade is more weakly structured and less tightly concentrated along a single diagonal. Although both one-dimensional marginals indicate where respondents are concentrated on each individual policy, they do not reveal how those respondents are paired across the two dimensions. Here, the heatmap provides information beyond the marginal densities by showing the joint arrangement of the mass in two-dimensional policy space.

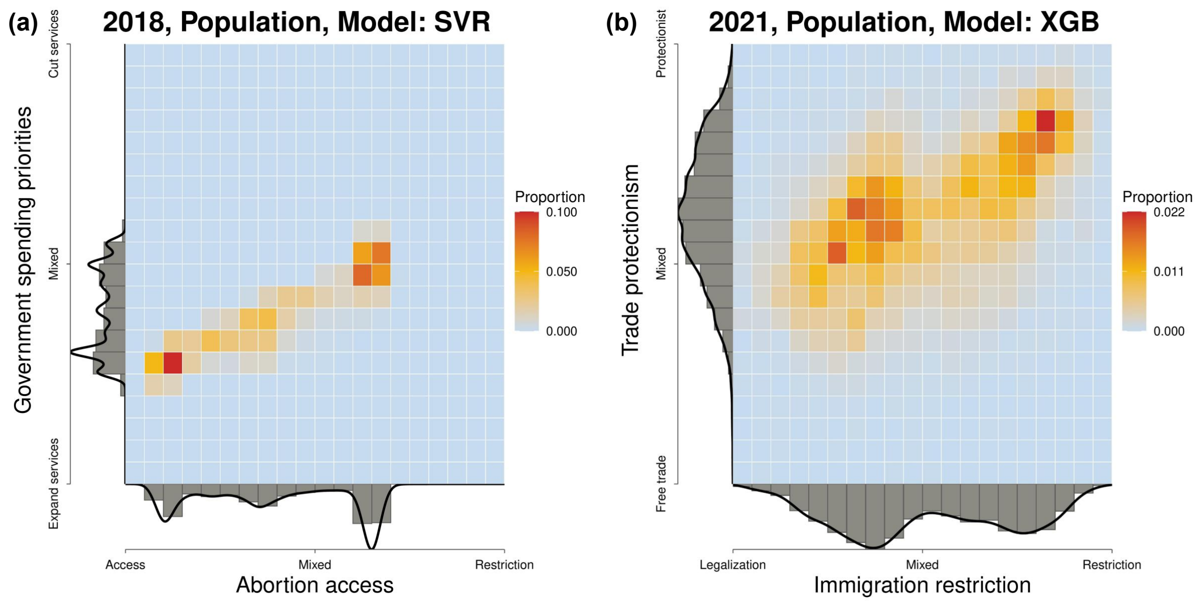

Figure 4. Contrasting joint structures in policy combinations. Comparison of two model-based heatmaps with different joint structures. Panel (a) illustrates a strong positive association, with the mass concentrated along a narrow diagonal band. Panel (b) illustrates a weaker association, where the joint distribution is more diffuse and the heatmap reveals structural features that are not apparent from the marginal density plots alone.

This comparison illustrates an important point: when two policy dimensions are highly correlated, the heatmap is visually simpler and closer to what the marginals imply. When the correlation is lower, the heatmap becomes more informative because it reveals joint structure that cannot be recovered from the one-dimensional density plots alone.

3.4 Resolving multipolar structure

Figure 5 shows the RF-based joint distribution of estimated preferences on public healthcare and abortion access in 2021, separately for the full population and for Republican, Independent, and Democratic respondents. Both policy scores are normalized to the interval $[0, 1]$. Higher healthcare scores correspond to stronger support for government-led healthcare, while higher abortion scores correspond to more restrictive attitudes toward abortion access. Because the plot is based on random forest estimates, it summarizes policy preferences after incorporating demographic covariates and their interactions.

The population-level distribution displays a clear multipolar structure rather than a single central cluster. One concentration appears near the market-based healthcare and restrictive-abortion region, while another appears closer to the government-led healthcare and abortion-access region. This indicates that the joint distribution of policy preferences is not well described by a single left–right average or by either policy dimension alone.

The partisan decomposition shows that partisanship is a major source of this structure. Republican respondents are concentrated closer to the market-based healthcare and restrictive-abortion region, while Democratic respondents are concentrated closer to the government-led healthcare and abortion-access region. Independent respondents occupy a more intermediate region, helping connect the two partisan poles in the aggregate distribution.

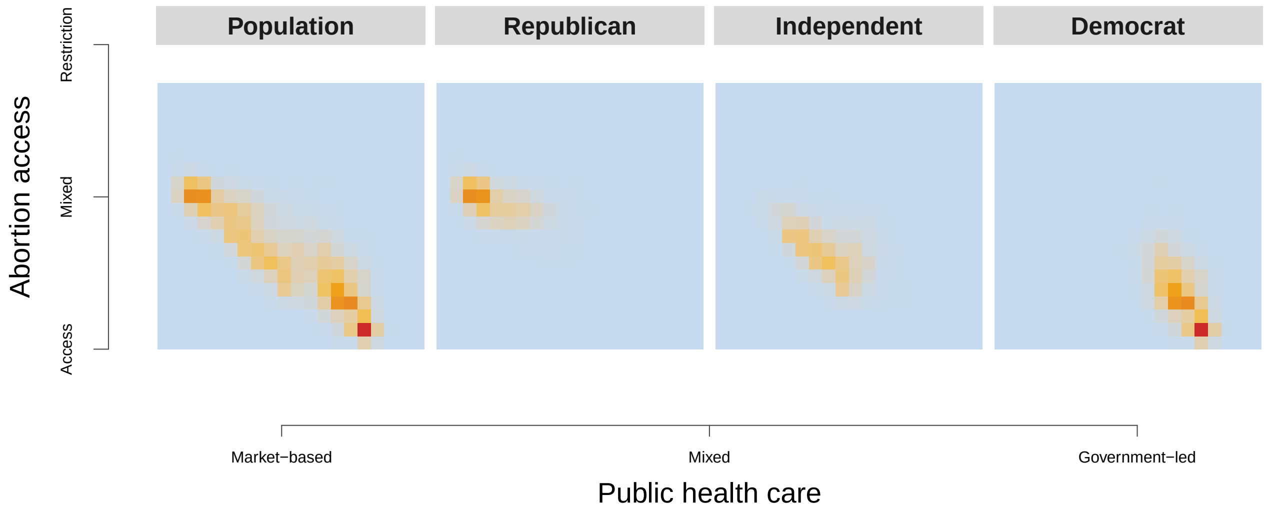

Figure 5. Decomposition of policy preference distribution by partisanship. RF-based heatmaps of estimated public healthcare preferences and abortion access preferences in 2021. The horizontal axis represents healthcare preferences, ranging from market-based to government-led, and the vertical axis represents abortion preferences, ranging from access to restriction. Columns show the full population, Republican respondents, Independent respondents, and Democratic respondents. Warmer colors indicate higher concentrations of respondents in the corresponding two-dimensional preference bin. The panels show how the aggregate distribution decomposes into distinct partisan-specific patterns.

More broadly, this two-dimensional decomposition illustrates the value of examining policy preferences jointly. A one-dimensional summary could show partisan differences on abortion or healthcare separately, but the heatmap shows how these preferences align within respondents and how the aggregate multipolar pattern arises from distinct partisan-specific distributions.

3.5 Capturing temporal dynamics of policy preferences

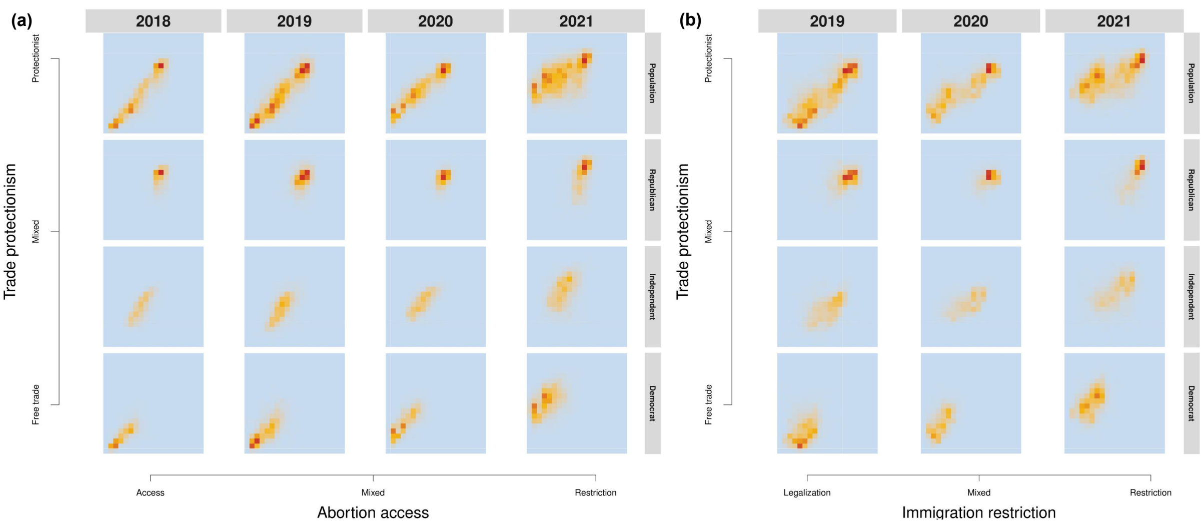

Figure 6. Temporal change in joint policy structures. Time-varying heatmaps showing the joint evolution of (a) abortion vs. trade and (b) immigration vs. trade preferences. Over time, the clusters shift across the two-dimensional policy space, reflecting changes in average attitudes. At the same time, the joint distributions become more diffuse, indicating weakening associations and increasing heterogeneity in policy alignment.

Figure 6 presents the joint distribution of individuals' policy preferences combining (a) abortion and trade and (b) immigration and trade, with trade scores interpreted such that higher values reflect more stringent foreign trade policies. Each panel shows population-level and partisan-specific distributions across multiple survey years.

1. **Shifts in the location of clusters on the policy plane:** Over time, the population distributions move across the two-dimensional space, indicating that the average relationship between abortion (or immigration) and trade preferences evolves. For instance, in the abortion-trade plots (Panel a), clusters shift upward along the trade dimension, suggesting a general trend toward stricter trade positions in later years. A similar shift is visible in the immigration-trade plots (Panel b), though with more variation across parties. These movements highlight how changes in broader political discourse or external events (e.g., global trade tensions, immigration debates) are reflected in joint policy attitudes.
2. **Increasing dispersion of the distributions:** In both abortion-trade and immigration-trade combinations, the distributions become more scattered over time. Earlier years show concentrated clusters, while later years display wider spreads across the 2D plane. This indicates a weakening of the correlation structure: abortion and trade, or immigration and trade, were once more tightly linked but now exhibit greater heterogeneity. Individuals are increasingly likely to hold stringent views on one issue while adopting moderate or permissive positions on the other, producing more cross-pressured or non-aligned combinations.

Together, these dynamics suggest that the electorate is not static in its policy bundling. Partisan alignment remains important, but both the center of mass (average position) and the strength of correlation between issues shift over time, reflecting broader political and social changes. Crucially, these temporal trends are only visible in the two-dimensional heatmaps—one-dimensional distributions of abortion, immigration, or trade alone would show marginal shifts but would not reveal the changing alignment between issue preferences.

3.6 Revealing subgroup variations and within-party structure

Figure 7 visualizes the joint distribution of Republican respondents' estimated scores on environmental protection (x-axis; higher = stronger pro-environment stance) and gun control (y-axis; higher = stricter regulation). The first row shows all Republicans, followed by subgroups defined by geography (Big Metro vs. Other Counties), education (College vs. Non-college), and gender (Female vs. Male).

1. **Aggregate Pattern:** Republicans as a whole exhibit a bimodal distribution with two clear poles: one group favoring both environmental protection and gun control, and another opposing both. This pattern reveals a significant within-party divide on how these two policy domains align.
2. **Geographic and Educational Subgroups:** The bimodal pattern persists when Republicans are divided by geography or education, suggesting that urban-rural differences and educational attainment do not explain the observed heterogeneity. Both subgroups maintain clusters at similar positions, indicating consistent polarization across these dimensions.
3. **Gender Subgroups:** When divided by gender, the two-pole structure largely disappears. Female Republicans cluster toward the pro-environment and pro gun-control region, whereas male Republicans concentrate near the opposite end. This finding suggests that gender differences are particularly visible within the Republican subgroup for this policy combination and are associated with the observed subgroup heterogeneity.

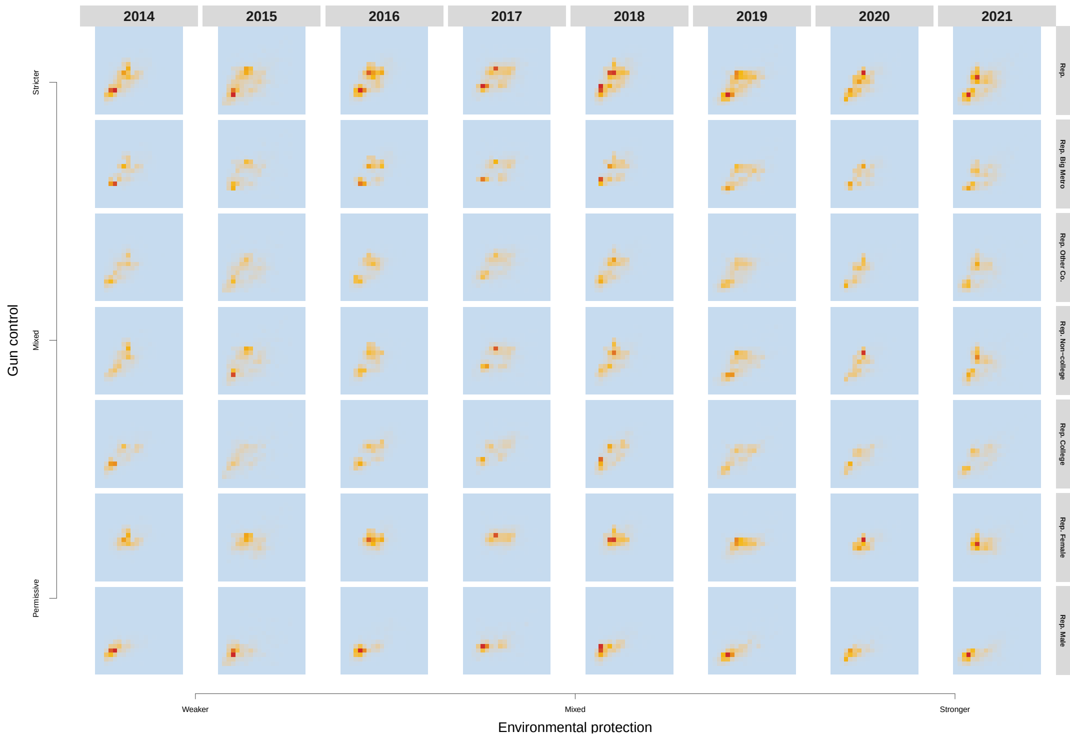

Figure 7. Within-party decomposition by demographic variables. Subgroup analysis of Republican respondents' joint preferences on environment (x-axis) and gun control (y-axis). The overall two-pole structure—pro-environment/pro-gun-control versus anti-environment/permissive-gun—remains visible after stratifying by geography and education, but largely disappears after stratifying by gender. This pattern suggests that gender accounts for a significant portion of the within-party heterogeneity in this joint policy space.

The two-dimensional visualization reveals ideological variation that conventional one-dimensional visualizations may overlook. By leveraging model-based estimated scores, this approach exposes demographically driven substructures—in this case, gender-based cleavages—that contribute to apparent within-party polarization. The framework thus provides a useful tool for exploring intra-party dynamics, identifying demographic sources of heterogeneity, and informing future research and policy analysis.

3.7 Interpreting subgroup heatmaps based on fitted scores

The subgroup heatmaps in Figures 5 and 7 are based on model-fitted policy scores and should therefore be interpreted as model-based subgroup summaries rather than purely descriptive raw-score distributions. Because subgrouping variables, such as partisanship and gender, are also included as predictors in the fitted-score models, subgroup differences may reflect both observed variation in policy responses and the model's representation of covariate-specific structure. In this sense, the visualization inherits the assumptions, smoothness, and flexibility of the fitted model. For example, heatmaps based on a linear model may primarily reflect additive and approximately linear covariate relationships, while more flexible specifications can represent localized nonlinearities and interactions more readily.

This model dependence is not a limitation unique to the heatmap display, but a general feature of any visualization based on fitted values. In our main subgroup examples, we use random forests because they provide a flexible nonparametric specification that can capture nonlinear and interaction patterns without requiring them to be specified in advance. Nevertheless, the resulting heatmaps should still be interpreted as smoothed model-based summaries, not as direct replacements for the observed score distributions. To aid interpretation, we provide corresponding raw-score subgroup heatmaps in the Supplementary Material (Figures S1–S2). Overall, the main subgroup patterns remain visible in the raw-score heatmaps, although the patterns are less visually distinct when subgroup combinations become more complex, as in the decomposition shown in Figure S2. This comparison clarifies that fitted-score heatmaps complement, rather than replace, raw-score visualizations.

4 DISCUSSION

The proposed framework is intended to provide an interpretable visualization of multidimensional policy-preference patterns across domains, years, and subgroups. A design choice in this framework is the construction of domain-level policy scores from multiple survey items. These scores should be interpreted as transparent descriptive indices rather than as full measurement models. Because the component items within each policy domain are weighted equally, each item contributes the same amount to respondents' locations in the two-dimensional policy space. This equal-weighting assumption supports reproducibility and consistent comparison across years, policy domains, and fitted-score models, but it may also affect the resulting heatmaps. Alternative weighting schemes or latent-variable approaches could place greater emphasis on particular items and may therefore shift some respondents' positions. Future work could extend the proposed visualization framework by incorporating alternative score-construction strategies and examining how sensitive the visualized policy-preference patterns are to those choices.

The proposed visualization framework demonstrates how integrating estimation models with two-dimensional heatmaps can enhance the analysis of complex survey data. Unlike conventional plots based on raw scores, the model-based approach generates visual representations that reflect how policy attitudes co-vary after incorporating multiple demographic and contextual factors. By allowing the estimation step to be implemented using linear models (LM), random forests (RF), support vector regression (SVR), or extreme gradient boosting (XGB), the framework is not tied to any single modeling choice. Instead, it provides a general visualization strategy that can be paired with a wide range of estimation methods.

By combining model-based estimation with interpretable visualization, the framework moves beyond traditional descriptive summaries to provide a model-informed view of survey data. The resulting heatmaps reveal joint policy structures, subgroup variation, and temporal dynamics within a common graphical framework, while the use of different estimation models demonstrates that the visualization approach is adaptable across both classical and machine-learning methods. This flexibility makes the framework a scalable tool for studying how multiple interdependent factors are associated with public attitudes.

Substantively, the empirical results show that the framework can uncover patterns that are difficult to detect from separate one-dimensional policy summaries. In particular, the heatmaps reveal multipolar structures in aggregate opinion, changes in the alignment between policy domains over time, and within-party subgroup differences that point to specific demographic sources of heterogeneity. These findings illustrate that the approach is not only a visualization tool, but also a useful way to examine how policy attitudes are organized, aligned, and differentiated across groups and years.

All analysis and visualization steps are implemented in the `surveyMDV` R package, which is available on GitHub (https://github.com/yu-jingtian/surveyMDV). The package provides functions for constructing and visualizing multidimensional policy-preference heatmaps, along with example data and worked examples demonstrating the complete workflow. These examples allow users to reproduce the figures and analyses presented in this paper and adapt the framework to their own survey datasets. The package is fully open source, enabling users to modify the functions, extend the visualization tools, and incorporate them into new applications. To improve figure accessibility, the accompanying package also provides an optional color-vision-accessible sequential palette through `palette = "accessible"`; an example is shown in Supplementary Material Figure S3.

When applied in policy-relevant settings, the proposed heatmaps should be interpreted as exploratory, observational, and hypothesis-generating visual summaries rather than as stand-alone bases for policy decisions or causal claims. The visualized patterns depend on the survey design, representativeness of the sample, construction of the policy scores, and the fitted models used to generate predicted scores. In particular, model-based heatmaps summarize fitted-score patterns and should be distinguished from raw observed public preferences. Because the framework describes associations rather than causal effects, applied interpretations should consider model uncertainty, survey weighting and coverage, and the substantive meaning of the underlying score construction before drawing policy conclusions.

The visualization framework can be readily applied to other large-scale survey datasets that capture multiple policy attitudes and demographic variables. Suitable examples include the American National Election Studies (ANES), the General Social Survey (GSS), the Pew Research Center's Political Typology Survey, additional waves of the Cooperative Election Study (CES), and international surveys such as the European Social Survey (ESS) and the World Values Survey (WVS). Beyond political opinion research, the same approach can be applied to issue-specific studies in health, environmental, and education policy. Future extensions may incorporate uncertainty visualization, dynamic or longitudinal displays, and alternative estimation methods to further enhance interpretability and robustness.

CONFLICT OF INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

Conceptualization: Jingtian Yu; Yanming Di; Sharmodeep Bhattacharyya. Methodology: Jingtian Yu; Yanming Di. Data curation: Jingtian Yu. Formal analysis: Jingtian Yu. Validation: Helen Chang; Neil Hwang. Writing - original draft: Jingtian Yu. Writing - review & editing: Helen Chang; Yanming Di; Sharmodeep Bhattacharyya; Shirshendu Chatterjee; Neil Hwang. All authors approved the final submitted draft.

FUNDING

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

DATA AVAILABILITY STATEMENT

The CES data utilized in this study are publicly available from the Harvard Dataverse: Cumulative CES Policy Preferences (V3) includes policy preference measures derived from the CES. It is available at: `https://doi.org/10.7910/DVN/OSXDQO`. Cumulative CES Common Content (V8) compiles responses from the CES common content surveys. It is available at: `https://doi.org/10.7910/DVN/II2DB6`.

The 2023 Rural-Urban Continuum Codes were obtained from the U.S. Department of Agriculture, Economic Research Service. The dataset is publicly available at: `https://www.ers.usda.gov/data-products/rural-urban-continuum-codes/`.

REFERENCES

Abramowitz, A. I. and Saunders, K. L. (2008). Is polarization a myth? *The Journal of Politics* 70, 542–555. doi:10.1017/S0022381608080493

Ansolabehere, S. and Jones, P. E. (2010). Constituents' responses to congressional roll-call voting. *American Journal of Political Science* 54, 583–597

Ansolabehere, S., Rodden, J., and Snyder Jr, J. M. (2008). The strength of issues: Using multiple measures to gauge preference stability, ideological constraint, and issue voting. *American Political Science Review* 102, 215–232

Axelrod, R., Daymude, J. J., and Forrest, S. (2021). Preventing extreme polarization of political attitudes. *Proceedings of the National Academy of Sciences* 118, e2102139118

Brown, T. E., Jauregui, G. P., Mettler, S., and Rivera, M. (2025). A rural-urban political divide among whom? Race, ethnicity, and political behavior across place. *Politics, Groups, and Identities* 13, 229–242. doi:10.1080/21565503.2024.2328551

Curini, L. and Franzese, R. J. (eds.) (2020). *The SAGE Handbook of Research Methods in Political Science and International Relations* (SAGE Publications)

[Dataset] Dagonel, A. (2021). Cumulative CES Policy Preferences. `https://doi.org/10.7910/DVN/OSXDQO`. doi:10.7910/DVN/OSXDQO. Harvard Dataverse, Version V3

Engebretsen, M. and Kennedy, H. (eds.) (2020). *Data Visualization in Society* (Amsterdam University Press)

Fiorina, M. P. and Abrams, S. J. (2008). Political polarization in the American public. *Annual Review of Political Science* 11, 563–588. doi:10.1146/annurev.polisci.11.053106.153836

Hagen, L., Keller, T. E., Yerden, X., and Luna-Reyes, L. F. (2019). Open data visualizations and analytics as tools for policy-making. *Government Information Quarterly* 36, 101387

Hansen, M. A. and Dolan, K. (2025). Cross-cutting identities in american politics: Gender, party, and attitudes on gun reform. *American Politics Research* 53, 293–305

Hetherington, M. J., Long, M. T., and Rudolph, T. J. (2016). Revisiting the myth: New evidence of a polarized electorate. *Public Opinion Quarterly* 80, 321–350

Kleinfeld, R. (2023). Polarization, democracy, and political violence in the united states: What the research says

[Dataset] Kuriwaki, S. (2023). Cumulative CES Common Content. `https://doi.org/10.7910/DVN/II2DB6`. doi:10.7910/DVN/II2DB6. Harvard Dataverse, Version V8

Kuriwaki, S., Ansolabehere, S., Dagonel, A., and Yamauchi, S. (2024). The geography of racially polarized voting: Calibrating surveys at the district level. *American Political Science Review* 118, 922–939. doi:10.1017/S0003055423000436

Langford, K. (2023). *The 2023 polarization index: A holistic analysis of the state of U.S. political strife*. Report, The Pell Center for International Relations and Public Policy at Salve Regina University. `https://salve.edu/pell-center/projects-and-initiatives/unpacking-political-polarization`

Lindh, A. and Andersson, A. B. (2025). Social networks and distributive conflict: the class divide in social ties and attitudes to income inequality across 29 countries. *European Sociological Review* 41, 163–178

Mason, L. (2015). “I disrespectfully agree”: the differential effects of partisan sorting on social and issue polarization. *American Journal of Political Science* 59, 128–145

McCabe, K., Zhu, Y., Bajaj, S. S., and Martin, A. F. (2024). Increasing voter participation through health care–based voter registration. *JAMA Health Forum* 5, e241563. doi:10.1001/jamahealthforum.2024.1563

Mehlhaff, I. D. (2024). A group-based approach to measuring polarization. *American Political Science Review* 118, 1518–1526. doi:10.1017/S0003055423001041

Nash, K., Trott, V., and Allen, W. (2022). The politics of data visualisation and policy making. *Convergence* 28, 3–12

Pew Research Center (2014). *Political polarization in the American public*. Report, Pew Research Center. `https://www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public/`

Reid, J. C., Brown, S. J., and Blount-Hill, K.-L. (2025). Political ideology and american attitudes toward police reform: The case of the conflicted conservative. *Socius* 11, 23780231251361695

Segaran, T. and Hammerbacher, J. (eds.) (2009). *Beautiful Data: The Stories Behind Elegant Data Solutions* (O’Reilly Media)

Stewart, A. J., Plotkin, J. B., and McCarty, N. (2021). Inequality, identity, and partisanship: How redistribution can stem the tide of mass polarization. *Proceedings of the National Academy of Sciences* 118, e2102140118

Tartaglione, J. and de-Wit, L. (2025). How the manner in which data is visualized affects and corrects (mis)perceptions of political polarization. *British Journal of Social Psychology* 64, e12787. doi:10.1111/bjso.12787

U.S. Department of Agriculture, Economic Research Service (2023). *Rural-Urban Continuum Codes*. `https://www.ers.usda.gov/data-products/rural-urban-continuum-codes/`

Wang, X., Sirianni, A. D., Tang, S., Zheng, Z., and Fu, F. (2020). Public discourse and social network echo chambers driven by socio-cognitive biases. *Physical Review X* 10, 041042

Zigerell, L. J. (2023). Introducing political science students to data visualization strategies. *Journal of Political Science Education* 19, 1–15. doi:10.1080/15512169.2022.2099879