

Issue Opinions in the U.S. States ^{*}

Michael Auslen[†]

September 15, 2026

Abstract

Measures of public support for specific policies at the state level are essential to key questions in political science: How well-represented is the public? How does opinion evolve in response to changes in political environments? Yet such measures are available only for recent years, largely because the long-running surveys that could extend them rely on cluster sampling, which standard methods for subnational opinion estimation handle poorly. I leverage recent advances in small-area estimation—Clustered Multilevel Regression and Poststratification (CMRP)—to construct a dataset of state-level opinion on 289 policy questions from the American National Election Study, covering 1952 to 2024. I apply these estimates to two questions. First, I recover the established finding that state policies are responsive to public opinion, particular from the 1970s onward. Second, I extend the study of nationalization to individual issues, showing that state publics’ preferences have largely converged over time, though the trend on racial issues is starting to reverse.

Word count: 6,665

Keywords: Public opinion; State politics; Multilevel Regression and Poststratification; Responsiveness; Nationalization

^{*}I thank Max Goplerud, Chenoa Jorgason, Isaac Mehlhaff, seminar participants at the Texas Junior Methods Workshop, and attendees at SPPC 2025, MPSA 2025, TexMeth 2026, and PolMeth 2026 for helpful feedback.

AI Disclosure: Claude (Sonnet 4.5) was used to assist in writing code to clean the large amount of survey data in this project; details are discussed in Appendix A. The text is human-authored; however, some text was provided to Claude (Opus 4.8) to assist with editing and organizational suggestions, some of which were implemented manually.

[†]Assistant Professor, University of Texas at Austin, Austin, TX. Email: mauslen@austin.utexas.edu. URL: <https://michaelauslen.com>.

Measuring what the public wants from government is integral to the study and practice of democratic representation. It is impossible to evaluate whether and when government takes actions that constituents support without having some measure of opinion at the appropriate geographic level of aggregation. For that reason, state-level public opinion about policy issues has proven especially valuable for studying representation in the states (e.g., Erikson, Wright and McIver 1993; Lax and Phillips 2012; Caughey and Warshaw 2018) and the U.S. Senate (e.g., Wright 1989; Fowler 2005; Lax, Phillips and Zelizer 2019). Likewise, data about public opinion is central to understanding how the views of the electorate have changed over time (e.g., Page and Shapiro 1992; Stimson, Mackuen and Erikson 1995; Caughey and Warshaw 2022).

Modern advances in opinion estimation have facilitated the measurement of policy preferences in the states, as well as at other subnational geographic levels. In doing so, scholars generally follow one of two conceptual approaches: summarizing policy preferences on a latent scale, such as ideology or policy mood, or estimating support for individual policies at the relevant levels of geography. Both are well-developed for measuring opinion in recent years, aided by large-N surveys such as the Cooperative Election Study (CES; formerly CCES). However, going back historically proves more challenging. While a number of scholars have produced valuable datasets of latent measures at the state level (Caughey and Warshaw 2015; Tausanovitch and Warshaw 2013; Enns and Koch 2013), no comparable data exists for issue-specific public opinion in the states.

A major reason for this is that workhorse surveys of American politics have long relied on sampling designs that modern methods for subnational opinion estimation handle poorly. In recent decades, Multilevel Regression and Poststratification (MRP; Park, Gelman and Bafumi 2004) has become the standard method for estimating subnational opinion from surveys conducted at higher geographic levels (Lax and Phillips 2009b; Warshaw and Rodden 2012). However, MRP estimates have higher-than-normal error when the surveys they use were constructed using cluster sampling (also called area probability sampling), as the ANES

has been for nearly all of its history (Stollwerk 2017). This occurs because cluster samples use a limited set of geographic areas to stand in for many similar places, producing samples that are nationally representative but not necessarily representative within smaller units like states. As a result, the state-level effects that MRP models estimate can be biased. The development of Clustered MRP (CMRP; Auslen 2024) provides a solution: By modeling opinion within the strata from which survey clusters are drawn, it recovers more accurate state estimates, even from surveys that were not designed to be representative at the state level.

This paper addresses this data problem by providing a novel dataset of state-level public opinion, spanning the full time series of the American National Election Study (ANES). Using CMRP to measure state opinion, it covers 289 distinct policy questions asked in the ANES over the period 1952–2024, each of which appears in an average of 8.8 survey years. The dataset is publicly available, including state-level estimates of opinion for all state-issue-years in the ANES, as well as standard errors of these estimates.¹ Constructing these estimates required substantial data-reconstruction work. Because the ANES sampling design has changed frequently over seven decades and its documentation varies in detail, I collected sampling reports for each survey and carefully matched the geographies in each design to the sampling procedures used, in order to build the poststratification frames that CMRP requires.

These estimates fill an important hole in the availability of public opinion data by offering a long-running time series of opinion on individual issues. They may be useful to scholars interested in a wide range of questions. In this paper, I apply them to two: First, measures of public opinion are essential building blocks of the study of political representation. Using the public opinion estimates developed here and an archive of state policy data (Grossmann, Jordan and McCrain 2021), I recover the standard finding that public opinion predicts policy decisions and extend it to the long-running span of the ANES. This long temporal coverage

¹Data are available for download on Dataverse at [URL TO COME].

allows me to demonstrate that the degree of responsiveness and policy congruence has changed over time. While there was essentially zero relationship between opinion and policy in the period 1952–1970, the connection grew rapidly until the early 1990s, before diminishing in recent years.

Second, I study whether and when public opinion on individual issues has nationalized. While American politics is widely understood to have nationalized over the 20th and early 21st century (Hopkins 2018), far less is known about the nationalization of public opinion on policy. I find that issue-specific opinion has largely converged across the states over the last 72 years, particularly through the 1970s. However, this pattern varies by issue domain, with attitudes on racial issues, in particular, beginning to show divergence across the states in recent years.

The paper proceeds as follows. First, I briefly review the existing approaches to measuring subnational opinion. Next, Section 2 describes both the ANES survey data and the CMRP model I use to estimate state opinion, including how it differs from traditional MRP approaches. In Section 3, I then validate the accuracy of the estimates against state election returns. I show that, using responses from the ANES, CMRP recovers ground truth election results better than alternatives (additional validation is demonstrated in the online Appendix). Finally, in Section 4, I apply the state estimates to two important substantive questions in the literature: responsiveness in state policymaking and the nationalization of public opinion.

1 Issue Opinions at Subnational Geographies

Scholars interested in measuring opinion in subnational geographies face an important choice: whether to summarize preferences on a latent scale, or to estimate them issue-by-issue. The first approach “lumps” individual policy questions together as components of a one- or multiple-dimensional latent trait (e.g., ideology, policy mood).² These latent measures can be

²Druckman and Jacobs (2006) coined the “lumping” and “splitting” terminology that is useful here for describing alternative approaches.

useful for describing and understanding broad shifts in the ideological leanings of the public (e.g., Ansolabehere, Rodden and Snyder 2008; Caughey and Warshaw 2022). Alternatively, scholars may “split” opinion into discrete issues, measuring mass support and opposition to particular policies. This can be especially useful when citizens hold idiosyncratic views that do not align to the typical left-right dimension of American politics (e.g., Broockman 2016; Broockman and Lauderdale 2025), and for research questions that rely on measures of domain-specific public opinion, such as responsiveness where it is important to measure opinion and policy action on the same scale (Achen 1978; Matsusaka 2001).

Regardless of the construct of interest, measuring public opinion at subnational levels can be difficult: Large amounts of survey data must be collected, and statistical models must be carefully designed to summarize opinion at relevant geographies with the smallest possible measurement error. For studies of contemporary public opinion (2000 and later), these challenges are greatly simplified by the advent of large-N surveys, such as the Cooperative Election Study (CES; formerly CCES), National Annenberg Election Survey (NAES), the Democracy Fund Nationscape study, and the CHIP50 study. These surveys aid “lumping” and “splitting” alike by making large samples freely available to be summarized using Item Response Theory (IRT) models (e.g., Caughey and Warshaw 2015; Tausanovitch and Warshaw 2013) and disaggregation. Likewise, advances in opinion estimation with Multilevel Regression and Poststratification (MRP; Park, Gelman and Bafumi 2004) have made subnational opinion estimation possible with much smaller surveys using off-the-shelf software (see fuller discussion in Section 2 below).

Constructing estimates of historical public opinion remains difficult. A number of scholars have published useful datasets of scaled opinion measures dating back to as early as 1936. In their influential study, Erikson, Wright and McIver (1993) used disaggregation to estimate a single snapshot of state ideology from surveys over the period 1976–88. More recently, Caughey and Warshaw (2018) used IRT models to estimate ideology on social and economic issues. Table 1 summarizes many of the efforts to publish state-level estimates of both scaled

Table 1: Sample of State-Level Policy Public Opinion Datasets

	Lumping	Splitting
Contemporary (2000s onward)	Tausanovitch and Warshaw (2013); Enns and Koch (2013); Carsey and Harden (2010); IRT or other scaling method on large-N studies (e.g., CES)	MRP on large-N studies or pooled samples (e.g., Lax and Phillips 2009 <i>a</i> , 2012; Warshaw and Rodden 2012)
Historical (1950s–1990s)	Caughey and Warshaw (2015, 2022); Erikson, Wright and McIver (1993); Enns and Koch (2013); Norrander (2001)	Limited to studies without cluster-sampling adjustment covering a limited set of issues: Shirley and Gelman (2015); Pacheco (2014); Brace et al. (2002); Johnson, Brace and Arceneaux (2005)

ideology and individual opinion measures, in both the historical and contemporary periods.³

However, there has been no systematic effort to capture state-level opinion on individual policy issues going back historically. As Table 1 illustrates, while scaled measures of public opinion extend back decades, the historical estimation of *issue-specific* opinion has been undertaken only in single-issue studies or those with few issues (e.g., Brace et al. 2002; Pacheco 2014; Shirley and Gelman 2015). No dataset covers the broad range of salient policies before the 2000s.⁴ A primary reason for this omission is that common methods for subnational opinion estimation (i.e., disaggregation and conventional MRP) produce substantially higher error when applied to polls built on cluster sampling (Stollwerk 2017), which was ubiquitous into the 1990s and is still used by the ANES today. Recent advances in small-area estimation relax this constraint, making comprehensive historical issue estimation feasible for the first time—via an approach I develop in Section 2.

³I exclude from Table 1 datasets constructed from sources other than surveys, e.g., Berry et al. (1998), who use congressional roll-call scores to estimate ideology.

⁴At the national level, there have been significant efforts to chart the over-time development of public opinion (e.g., Page and Shapiro 1992) and to gather archives of survey results (e.g., the Roper Center for Public Opinion Research).

2 Data and Estimation

Measuring historical public opinion requires long-running sources of data. To that end, this paper uses the ANES to construct a dataset of Americans’ preferences over specific policy debates at the state level. The ANES is among the most consequential surveys for understanding the dynamics of opinion in the United States. Because of its high frequency and longevity—it has been fielded every two-to-four years since 1948—the ANES captures Americans’ changing views on the issues of the day. However, for nearly all of its history, the ANES has drawn respondents using cluster sampling (also called area-probability sampling). This made face-to-face interviewing more efficient but raises challenges for standard small-area estimation methods. This section describes the ANES data I use to construct the historical opinion dataset and the estimation approach I adopt to address these problems.

2.1 Policy Opinion in the ANES

To measure state-level opinion, I first collect a comprehensive dataset of individual responses to all policy-related questions in the ANES from 1952–2024.⁵ This includes all questions focused on the public’s preferences over government policy: those that ask about specific issues (e.g., ideal policies on abortion or levels of taxation), government spending (e.g., whether education funding should increase), and the role of government (e.g., whether the federal government should be stronger or weaker or whether government ought to protect civil rights). I exclude questions focused on evaluating candidates, parties, groups, and the efficacy of government decisions.

In all, I collect 289 unique survey questions, each of which is asked on an average of 8.8 surveys (the average survey includes 31.9 questions, though this number varies over time). Where questions ask for preferences on a scale, I binarize them by grouping “strongly agree” with “somewhat agree” and for Likert scales by combining responses above the midpoint.

⁵I exclude the 1948 survey from the dataset because it lacks geographic identifiers (i.e., state), rendering the MRP methods I employ unreliable.

This is common practice in studies using MRP, which almost always estimates opinion using binary variables. I discuss these coding decisions further in Appendix A; Appendix H lists the included questions.

2.2 Estimating Opinion from Cluster-Sampled Polls

In recent decades, Multilevel Regression and Poststratification (MRP) has become a workhorse method for estimating public opinion at subnational geographies (Gelman and Little 1997; Park, Gelman and Bafumi 2004). By modeling survey responses with demographic and geographic predictors, MRP pools information about respondents across geographies, allowing researchers to estimate support for policies in the states (or smaller geographies) from national surveys with relatively small state-level samples (Lax and Phillips 2009*b*; Warshaw and Rodden 2012; Tausanovitch and Warshaw 2013). For brevity, I describe the typical MRP approach in Appendix B.

However, MRP’s usefulness for historical public opinion has been limited by those surveys’ reliance on cluster sampling. These survey designs draw a set of cities or counties (called primary sampling units) to stand in for other, similar places and then randomly sample progressively smaller geographies until households are reached (Kish 1965). This approximates a random national or regional sample for face-to-face interviews at low cost, which explains its popularity with early survey researchers (Berinsky 2017). However, it complicates MRP because respondents are not drawn at random *within* states (Stollwerk 2017). As a result, the state-level geographic effect in a typical MRP model conflates the true state effect on opinion with features of the sampling frame itself. For example, a sample might include only a single cluster from Wisconsin, say urban Milwaukee, which stands in for other urban Midwesterners. This cluster is surely not representative of all Wisconsinites, in part because the rural parts of the state may be represented in the sample by rural clusters in Minnesota or Michigan.

Clustered MRP (CMRP) addresses this problem by bringing information from the sampling

frame into the opinion model (Auslen 2024).⁶ The key insight is that although a cluster-sampled poll is not representative of any individual state, it is designed to be representative of particular strata (e.g., the urban, suburban, and rural groupings from which clusters are drawn). CMRP therefore models opinion at the level of the state-by-stratum and then aggregates up to the state level by reweighting strata to their true shares of each state’s population. This also allows the core advantage of MRP to be retained, as demographic information is pooled across strata, regions, and states. This method produces lower levels of error compared to traditional MRP when used with cluster-sampled polls.

2.3 CMRP and the ANES

Before CMRP can be applied to the ANES, the strata that define each survey’s sampling frame must be reconstructed. Recall that CMRP models opinion at the state-by-stratum level and then aggregates to the state. This requires assigning each survey respondent to a stratum, and, for the poststratification stage, determining the share of each state’s population that falls in each stratum.⁷ Both depend on the same stratum definitions, which I recover from the sampling details often included with survey documentation.

However, documentation varies widely in its detail and consistency. Some specify the exact inclusion criteria for sampling strata and the specific steps in the sampling procedure, while others describe the frame only in broad terms. The particulars of these sampling frames changed frequently over the seven decades of the survey. Reconstructing each survey’s strata required assembling these reports individually, reconciling their definitions, and hand-matching counties and metropolitan areas to the Census units from which I build poststratification frames. Because metropolitan-area and urban classifications themselves changed repeatedly over this period, I match each survey to the geographic definitions in force at the time

⁶Caughey and Warshaw (2022) offer a partial solution with similar inspiration, modeling their group-level IRTs with a variable for urban/rural. This addresses a concern of cluster-sampled polls in practice though does not wholly incorporate the sampling frame.

⁷For ANES years 2004–2016, respondents’ sub-state geographic information (i.e., metropolitan area or county) is only available in restricted-use data. For all other years, geographic information necessary to reconstruct strata is in publicly available files.

rather than applying a single crosswalk. For some years, this required consulting external government administrative data or drawing on detailed technical reports published by the ANES and the Survey Research Center (e.g., Kish and Hess 1965; Heeringa, Connor and Darrah 1986).

Commonly, the ANES constructed its sampling frame by ranking metropolitan areas by population and grouping them into tiers accordingly. The largest areas (e.g., the top six or 12) were usually included with certainty, with one or two additional tiers of successively smaller metropolitan areas and a set of counties outside any metropolitan area comprising the rest of the sample. These tiers, crossed with the urbanness of geographies (central cities, suburbs, outlying areas, and rural places) form the strata for most ANES surveys. Because the particulars of the sampling frame differ across surveys, I reconstruct strata separately for each, assigning respondents to strata and computing each state-by-stratum population share from Census data. This is discussed in greater detail and for each survey year in Appendix C.⁸

Like standard MRP, CMRP proceeds in two stages. First, I fit a hierarchical model predicting individual binary responses y_i to each survey question. Because the sampling-frame strata change from one survey to the next and because pooling across years introduces additional assumptions about the stability of demographic and geographic effects, I model each survey-question-year separately rather than fitting a dynamic model (e.g., Mehlhaff 2024). I model opinion according to the following model, fit using the R package `vglm`, which recovers posterior means closely aligned with a fully Bayesian estimate (Goplerud 2024). I use unfactorized variational inference (which doesn't assume independence between random effect blocks), so the model also produces reliable uncertainty quantification (Goplerud,

⁸There are a number of variations and exceptions to this standard design. For example, in two years, 1978 and 1980, the sampling frame is built entirely on sets of congressional districts stratified by geographic and non-geographic similarity; some survey years include partial samples from random digit dialing or Web sampling; and the 2020 ANES was conducted entirely via random sample because the Covid-19 pandemic prevented in-person interviews. I detail the varying sampling frames and discuss how I model each in Appendix C.1.

Papaspiliopoulos and Zanella 2025), which I use to compute standard errors on the estimates.

$$\begin{aligned}
\Pr(y_i = 1) &= \text{logit}^{-1} \left(\beta_0 + \alpha_{s[i]}^{\text{state}} + \sum_{g \in G} \alpha_{j[i]}^g \right) \\
\alpha_s^{\text{state}} &\sim N \left(\alpha_{r[s]}^{\text{region}} + \sum_k f_k(X_{sk}), \sigma_{\text{state}}^2 \right) \\
\alpha_r^{\text{region}} &\sim N(0, \sigma_{\text{region}}^2); \quad \alpha_j^g \sim N(0, \sigma_g^2) \text{ for all } g \in G.
\end{aligned} \tag{1}$$

$\alpha_{s[i]}^{\text{state}}$ is a random intercept for state, which is modeled as a function of region and a number of continuous “contextual” covariates X_{sk} measured at the state level. These include Republican presidential vote share, plus % Union for economic and spending issues (Tausanovitch and Warshaw 2013) and % Evangelical for all other issues (Lax and Phillips 2012), modeled as splines $f(\cdot)$ to capture nonlinearity. The remaining $\alpha_{j[i]}^g$ terms are random intercepts for demographic factors (race, sex, education, and age) and the sampling frame stratum, as well as the two-way interactions among these to produce a “deep” model (see Ghitza and Gelman 2013).⁹ The interactions of stratum with state and region allow the model to include information about the structure of the sampling frame. The factors in G and the number of levels vary across survey years and with changing Census data; Appendix C details this variation.

Next, I poststratify. Using the hierarchical model, I predict the level of support for each question among all possible combinations of demographic and geographic attributes (e.g., the predicted response for Black women with college degrees between the ages of 18-29, living in the rural stratum of Alabama). These predicted values are then aggregated to the state level by taking a weighted average using the known distribution of the population, measured from Census data, as weights. I construct cell weights from Census data by assigning metropolitan areas or counties to the appropriate population-based tiers in the ANES sampling frame and use Census designations of central cities and urban and rural areas to define strata.

⁹Goplerud (2024) shows that “deep” hierarchical models that contain interactions perform comparably to machine-learning methods for MRP (e.g., Bisbee 2019; Ornstein 2020).

This procedure yields an estimate of support for each policy question in each state and year. I also compute standard errors for the opinion estimates by drawing 1,000 samples from the model’s variational approximation, poststratifying each, and taking the standard deviation of the resulting state estimates (see Appendix D for details). The resulting dataset is the core contribution of this paper. I validate these estimates in the following section.

3 Validation


To assess the accuracy of the estimates, I use CMRP to estimate state-level presidential vote share and compare these against the true election results. Vote choice is a common benchmark for validating opinion estimates, largely because it is one of the few attitudes for which a state-level ground truth is observed (e.g., Lax and Phillips 2009*b*; Warshaw and Rodden 2012). Auslen (2024) shows that CMRP out-performs traditional MRP and disaggregation when using cluster-sampled surveys. Here, I show that CMRP performs well when used with the ANES.

For each state in each presidential election year from 1952–2024, I fit a CMRP model of Republican vote share, as described in Equation 1 above, using all ANES respondents who report voting in the presidential election.¹⁰ I exclude the spline on presidential vote, as this is the benchmark against which estimates are compared. I then compare these estimates against the vote share received by the Republican in each state.

Figure 1 plots estimated Republican vote share against actual results for each state-year. The left panel shows estimates from disaggregation, while the right panel shows estimates from CMRP. CMRP estimates closely track the true results. The mean absolute error (MAE) of the estimates is 0.060, just slightly higher than Lax and Phillips (2012) report for similarly sized subsamples of a larger poll, constructed via a simple random draw. Disaggregation performs far worse (MAE = 0.111). Importantly, disaggregated estimates cannot be computed

¹⁰For elections prior to the ratification of the 26th Amendment in 1971, which lowered the voting age to 18, I limit the survey and poststratification data to people aged 21 or older.

Figure 1: Validation of CMRP Estimates with Election Results



Note: Points indicate observed and estimated Republican presidential vote share in a given state-year, using CMRP and disaggregation. For illustration, the 279 state-years that do not appear in the ANES (and so estimates cannot be produced via disaggregation) are plotted at zero. Green points (●) correspond to estimates for the South in the 1964 presidential election.

for 279 state-years in which the ANES sampled zero respondents. These state-years are reported at 0% on the y-axis of the figure for illustrative purposes, but are excluded from the error estimates. CMRP also out-performs traditional MRP using the ANES, which I show in Appendix E.

The largest errors are for the South during the 1964 election, reported in green in Figure 1. This is not surprising given the irregularity of that election: Barry Goldwater carried the Deep South, despite the region’s strong Democratic lean (Converse, Clausen and Miller 1965).¹¹

One tradeoff to using election results for validation is that while they present a rare opportunity to compare against a ground truth, vote choice variables may perform differently from the policy questions at the heart of this paper. In Appendix E, I compare CMRP estimates using the ANES with results for similar questions on two large-N studies covering the period 2000–2024, the NAES and CES. The appendix also validates that the state estimates from the models using `vglmr` are nearly identical to fully Bayesian estimates from

¹¹Other state opinion estimation efforts have encountered similar challenges (e.g., Lagodny et al. 2023).

Stan, and standard errors closely track them.


4 Applications

The dataset provides estimates of state opinion on 289 questions in all 50 states from 1952–2024, offering a broad view of how Americans’ policy preferences vary across space, time, and issues. Figure 2 illustrates this variation for a selection of nine issues in four states: California, Indiana, Massachusetts, and Texas. Point estimates report the share of the public that support the stated policy; error bars report the standard error of that point estimate. Because each year is modeled independently, no point estimates exist for years in which the question was not asked.

The figure points toward interesting variation both within and across states. First, looking across states, we can see that state-level preferences largely align with what we might expect. In general, states whose residents can be expected to hold more liberal positions on social issues generally do. For example, people in California and Massachusetts are more supportive of abortion rights than those in Indiana or Texas, while the latter are more supportive of the death penalty—especially in more recent years. On some issues, such as support for increased spending on public schools, opinion is highly correlated across states over time, while on other issues like health care, there is considerable variation. On some issues, notably isolationism and foreign trade, opinion across states is highly correlated and seems to follow broader national trends as public sentiment in the United States became increasingly, and later decreasingly, interested in globalization.

In the remainder of this section, I apply these estimates to two prominent debates in American politics. First, I present over-time evidence that state policy is responsive to public opinion. And second, I explore how the nationalization of public opinion has varied across issues and over time.

Figure 2: State Opinion over Time



Note: Plot shows state-level opinion estimates for four states on a selection of issues over the period 1952–2024. Error bars report standard errors of estimates.

4.1 Policy Responsiveness

A foundational question about representation in democracies is whether governments enact policies that the public wants. In the states, scholars have routinely found that policies are responsive to public opinion (Erikson, Wright and McIver 1993; Lax and Phillips 2009*a*; Caughey and Warshaw 2022), although these policies only align with the preference of the majority about half of the time (Lax and Phillips 2012). These studies, however, have typically been limited to recent decades or narrow snapshots in time, or have relied on scaled measures of ideology or mood to trace responsiveness over longer periods of time.¹²

I use the state opinion estimates from the ANES to study responsiveness across a six-decade time span. I collect policy and state expenditure data from the Correlates of State Policy Project (Grossmann, Jordan and McCrain 2021), and study whether states are more likely to enact policies when the public is more supportive of them. I focus on 22 specific policies that appear in at least two survey years and have a clear match to a substantive policy in the Correlates dataset.¹³ I likewise examine whether state spending in 8 specific areas is responsive to mass support for budget increases, again using issues included in both datasets. Appendix F lists the specific policies, as well as the original data sources collected in the Correlates data. The data are limited by state policy availability after 2016, so the analyses end in that year. I also take care to pair policies in year t with state opinion estimated in year $t - 1$, using the most recent estimate within the preceding six years if the ANES did not ask about a particular policy in $t - 1$.¹⁴


For the 22 policies, I model whether states have enacted the specific policy with a logistic regression on opinion, with policy-specific random intercepts and slopes, so that responsiveness may vary across policies. For the 8 spending areas, I model the two-year change in per-capita

¹²Caughey and Warshaw (2022), which primarily uses scaled ideology measures to study responsiveness over a long historical period, also presents results on individual issues using pooled surveys from the 1930s onward.

¹³In all, more than half of issues recur in at least two surveys, though many do not match to policies on the state agenda and included in the Correlates data.

¹⁴This ensures that opinion is always measured temporally prior to policy. Appendix F shows results are similar if only contemporaneous opinion is used.

Figure 3: Policy Responsiveness in the States



Note: Curves are predicted probabilities of specific policies being enacted (black solid line) and predicted spending change, standardized within domain (gray dashed line), modeled as a function of public opinion. 95% credible intervals are indicated with thinner curves. Full regression tables can be found in Table A4 in Appendix F.

spending—in 2000 dollars and standardized within each area to address variation in the magnitude of spending across categories—as a spline-smoothed function of opinion, again with area-specific random intercepts and slopes.

Figure 3 reports summary results, with the predicted probability curve from the policy analysis represented by the black solid line and the predicted change in standardized per-capita spending plotted as a gray dashed line. I find that as public support for policies increases, states are more likely to enact them. State budget changes are less affected by public opinion, though at the upper end of public support, the effect is larger. The data also allow me to compute the level of congruence between opinion and policies (the share of issue-state-years where policy or spending changes match the majority of the public). I find congruence on 52.4% of policies, consistent with prior work, and observe that the direction of inflation-adjusted per-capita spending changes matches public opinion in 55.2% of cases.

These results confirm the findings of prior work: when pooling over the full time period of the ANES, state governments are *responsive* to public opinion but are only *congruent* in

Table 2: Responsiveness on Specific Policies over Time

Period	Responsiveness	Congruence	N
1952–1970	-0.06 (2.01)	42.60%	899
1972–1990	9.92 (1.63)	56.98%	1,113
1992–2016	4.22 (2.14)	54.60%	2,548

Note: Table reports coefficients from responsiveness regressions (with standard errors in parentheses) and rates of congruence for policies and estimated opinion in the listed years. Full regression tables can be found in Table A5 in Appendix F.

about one-half of cases. I also uncover heterogeneity across issues: In particular, I see much stronger effects for high-salience social issues such as abortion (post-*Roe v. Wade*), LGBTQ rights, and race, and weaker effects on lower-salience issues. I discuss these in greater detail in Appendix F.

One advantage to this long-running set of opinion data is that it facilitates studying heterogeneity in representation over time. Table 2 examines whether responsiveness has changed over this long period, dividing it into three eras.¹⁵ The relationship between opinion and policy is strong and statistically significant from 1972–1990 but is not detectable in the earliest period (1952–1970), when congruence is also less than 50%. In more recent years, the relationship is considerably diminished. One challenge is that questions enter and exit the dataset, affecting the apparent relationship between opinion and policy (in particular, recent surveys include a larger number of policy questions). Still, these results suggest that state governments have become more responsive to public opinion over the course of the 20th century, though this has diminished in more recent years.

4.2 Nationalization of Issue Opinion

A prominent finding in recent scholarship is that American politics has nationalized. In particular, voter behavior is increasingly organized around national issues and national party identities, rather than distinct state or local ones (Hopkins 2018). While there are a number

¹⁵I use three periods of about 20 years each, rather than finer intervals because shorter windows contain too few policies to estimate responsiveness reliably.

of implications for this broad trend in American politics, the most well-studied among the public have been electoral: the increasing alignment of state and local election outcomes with national partisan tides (e.g., Sievert and McKee 2019; Rogers 2016) and declines in split-ticket voting (Jacobson 2015; Kuriwaki 2020). Some research has studied the nationalization of partisan ideologies within states (Caughey, Dunham and Warshaw 2018), but we know much less about whether the public’s preferences on particular *issues* have grown more alike across states. The state-level opinion estimates developed here make it possible to study the nationalization of issue opinion directly over a wide range of policies and a long period. I find that opinion nationalized over the mid-to-late 20th century, but that this trend has varied across issues.

Here, I study nationalization by measuring the dispersion of state-level public opinion on individual issues. For each issue-year, I compute $100 - IQR_{it}$, where IQR_{it} is the interquartile range of the state estimates for issue i in year t . Though a blunt measure, this has a natural interpretation: when state opinion has converged over a narrow range (i.e., when an issue has nationalized), $100 - IQR_{it}$ is higher, whereas when state publics are more distinct from one another, the measure is lower. This measure allows straightforward comparisons across issues and over time.¹⁶

I begin by asking which issues have most nationalized during the period covered by the ANES. I fit a hierarchical model of nationalization ($100 - IQR_{it}$) on time, with issue-specific random intercepts and slopes, and I rank the issues by these slopes. Larger positive slopes indicate issues on which opinion has most converged over time, whereas more negative slopes are those where opinion has most diverged. I restrict the analysis only to those issues that were included at least five times on the ANES, allowing for a long enough time span to identify a trend. This leaves 53 distinct issues.

Table 3 reports the issues that nationalized and diverged most during this period. The issues on which opinion most converged are largely those on which national consensus took

¹⁶In Appendix G, I show that my main results are generally robust to several alternative measures, as well as a model-based approach for estimating state-level opinion dispersion.

Table 3: Nationalization Trends by Issue

Most Nationalizing		Most Diverging	
Issue	Slope	Issue	Slope
Integrate Schools	0.48	Busing	-0.21
Integrate Neighborhoods	0.39	Equal Economic Opportunity	-0.19
School Prayer	0.38	Environmental Protection	-0.17
Desegregate Society	0.22	Death Penalty	-0.17
LGBT Antidiscrimination	0.22	Abortion: Ban	-0.17
Spending: Social Security	0.15	Affirmative Action	-0.15
Gays in Military	0.14	Spending: AIDS Crisis	-0.14
Spending: Unemployment	0.14	Foreign Aid	-0.13
Spending: Public Schools	0.11	Same-Sex Marriage	-0.13
Spending: Welfare	0.09	Vietnam: Pull Out	-0.12


Note: Issues sorted by slope from a hierarchical model on year. The Most Nationalizing issues are those where opinion has converged the most, and the Most Diverging issues are those whose trendlines have been least convergent. Issues included are those asked in at least five ANES surveys.

hold, especially in the earlier part of the period. In particular, questions related to integration and civil rights saw overall nationalization during the period, though, as I show below, opinion on more contemporary racial issues has diverged in recent years. Many of the most nationalizing issues are those on which prior work has found broad, nationwide opinion change (e.g., Page and Shapiro 1992). Conversely, I find the most divergent issues to be those that have become sources of widespread political conflict in recent years, such as abortion, LGBTQ rights, and the environment. Likewise, I find evidence of divergence for some issues that had relatively short but contentious lifespans in American politics (e.g., Vietnam and busing) but were removed from the ANES before opinion might have coalesced.

I next examine the trend of nationalization over time. I fit a hierarchical generalized additive model of $100 - IQR_{it}$ on year, with splines capturing the overall trend over time, domain-specific trends, and random effects for domain and issue that absorb time-invariant differences in levels of agreement across issues.¹⁷ As before, I restrict attention to issues asked at least five times. Figure 4 plots the estimated trend over 1952–2024, pooling over

¹⁷The model is described more fully in Appendix G.

Figure 4: Nationalization on Policy Issues




Note: Curve plots the trend of nationalization over time, estimated with a hierarchical generalized additive model, with the 95% confidence interval. Some outlier points are excluded to better elucidate trendlines. A table describing the full model can be found in Table A6 in Appendix G.

all issues. States began the period divided on the issues included in the ANES, before converging rapidly over three decades. Since then, nationalization has largely leveled off. State-by-state agreement remains higher today than in the 1950s and '60s. A key limitation is that I only observe opinion on questions that were included in the ANES, which may not be representative of all issues, though the random effects model helps to smooth out some of these differences.

Nationalization has also varied across policy domains. Figure 5 plots the trends by domain, estimated using the same model described above. Opinion has consistently been most nationalized on economic and foreign policy issues, while race and social/cultural issues have been more distinct across states. This may be unsurprising given the distinctiveness of the American South. Across issues, the global trend toward nationalization can be observed. This trend is most pronounced on racial and social/cultural issues, which converged rapidly through the 1970s. However, opinion on racial issues has begun to diverge again since around

Figure 5: Nationalization and Issue Convergence by Domain



Note: Estimated nationalization trends by policy domain, from a hierarchical GAM; higher values indicate greater cross-state agreement. A table describing the full model can be found in Table A6 in Appendix G.

2000.

Taken together, these results show that the nationalization of American politics extends to the realm of issue-specific opinion. Over the second half of the 20th century, state publics' opinions converged substantially, and they remain nationalized today. Underlying these trends is considerable heterogeneity across policy domains, especially on racial issues.

5 Conclusion

This paper has introduced state-level estimates of public opinion on all policy issues in the ANES over the period from 1952–2024. Covering 289 distinct policies, it is the first large-scale effort to produce publicly available, issue-specific opinion measures at subnational geographies for the full range of topics covered on the ANES. To do so, I relied on Clustered MRP, a recent development in subnational opinion measurement that improves estimation from cluster-sampled surveys compared to standard methods.

I illustrate the utility of these estimates by applying them to two important questions in the literatures on state politics and public opinion. First, I show that state policy is responsive to public opinion on individual issues, recovering a common finding in the literature using opinion estimated from the ANES. Interestingly, I find that this relationship was weak in the mid-20th century before strengthening considerably. Second, I examine whether opinion on individual issues has nationalized, extending our understanding of a key trend in contemporary politics beyond election results and ideologies. I find that state publics have grown more similar over the past seven decades, though the trends vary across issues—particularly racial issues, on which states are increasingly diverging.

In addition to the applications studied here, these estimates will be useful to scholars interested in a range of other questions. First, they may be useful in studying representation in other contexts (e.g., the U.S. Senate), and the downstream consequences of changes in public policy. Because many ANES questions recur over decades, they are also well suited to

studying the dynamics of public opinion in the states. To support future work, the full dataset is released publicly on Dataverse, including state-level estimates and their standard errors, as is the code needed to reproduce them. I also release the poststratification files constructed for use with the ANES, which required recovering the sampling design of each survey from historical technical reports and which can be reused for other small-area estimation tasks with the ANES.

More broadly, these estimates introduce historical depth, which has long been available only for scaled measures, to the study of individual issue opinion. In doing so, they allow scholars to trace public opinion on individual issues historically across the states and make it possible to better understand key questions about representation and the dynamics of American public opinion.

References

- Achen, Christopher H. 1978. “Measuring Representation.” *American Journal of Political Science* 22(3):475–510.
- Ansolabehere, Stephen, Jonathan Rodden and James M. Snyder. 2008. “The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting.” *American Political Science Review* 102(2):215–232.
- Auslen, Michael. 2024. “Improving Subnational Opinion Estimation from Cluster-Sampled Polls.” *State Politics & Policy Quarterly* 24(4):447–467.
- Berinsky, Adam J. 2017. “Measuring Public Opinion with Surveys.” *Annual Review of Political Science* 20(1):309–329.
- Berry, William D., Evan J. Ringquist, Richard C. Fording and Russell L. Hanson. 1998. “Measuring Citizen and Government Ideology in the American States, 1960-93.” *American Journal of Political Science* 42(1):327.

- Bisbee, James. 2019. "BARP: Improving Mister P Using Bayesian Additive Regression Trees." *American Political Science Review* 113(4):1060–1065.
- Brace, Paul, Kellie Sims-Butler, Kevin Arceneaux and Martin Johnson. 2002. "Public Opinion in the American States: New Perspectives Using National Survey Data." *American Journal of Political Science* 46(1):173–189.
- Broockman, David E. 2016. "Approaches to Studying Policy Representation." *Legislative Studies Quarterly* 41(1):181–215.
- Broockman, David E. and Benjamin E. Lauderdale. 2025. "'Moderates'." *American Political Science Review* pp. 1–10.
- Carsey, Thomas M. and Jeffrey J. Harden. 2010. "New Measures of Partisanship, Ideology, and Policy Mood in the American States." *State Politics & Policy Quarterly* 10(2):136–156.
- Caughey, Devin and Christopher Warshaw. 2015. "Dynamic Estimation of Latent Opinion Using a Hierarchical Group-Level IRT Model." *Political Analysis* 23(2):197–211.
- Caughey, Devin and Christopher Warshaw. 2018. "Policy Preferences and Policy Change: Dynamic Responsiveness in the American States, 1936–2014." *American Political Science Review* 112(2):249–266.
- Caughey, Devin and Christopher Warshaw. 2022. *Dynamic Democracy: Public Opinion, Elections, and Policymaking in the American States*. Chicago Studies in American Politics Chicago (Ill.): the University of Chicago press.
- Caughey, Devin, James Dunham and Christopher Warshaw. 2018. "The Ideological Nationalization of Partisan Subconstituencies in the American States." *Public Choice* 176(1):133–151.
- Converse, Philip E., Aage R. Clausen and Warren E. Miller. 1965. "Electoral Myth and Reality: The 1964 Election." *American Political Science Review* 59(2):321–336.
- Druckman, James N. and Lawrence R. Jacobs. 2006. "Lumpers and Splitters: The Public Opinion Information That Politicians Collect and Use." *Public Opinion Quarterly* 70(4):453–476.

- Enns, Peter K. and Julianna Koch. 2013. "Public Opinion in the U.S. States: 1956 to 2010." *State Politics & Policy Quarterly* 13(3):349–372.
- Erikson, Robert S., Gerald C. Wright and John P. McIver. 1993. *Statehouse Democracy: Public Opinion and Policy in the American States*. New York: Cambridge University Press.
- Fowler, James. 2005. "Dynamic Responsiveness in the U.S. Senate." *American Journal of Political Science* 49(2):299–312.
- Gelman, Andrew and Thomas C Little. 1997. "Poststratification Into Many Categories Using Hierarchical Logistic Regression." *Survey Methodology* 23(2):127–35.
- Ghitza, Yair and Andrew Gelman. 2013. "Deep Interactions with MRP: Election Turnout and Voting Patterns Among Small Electoral Subgroups." *American Journal of Political Science* 57(3):762–776.
- Goplerud, Max. 2024. "Re-Evaluating Machine Learning for MRP Given the Comparable Performance of (Deep) Hierarchical Models." *American Political Science Review* 118(1):529–536.
- Goplerud, Max, Omiros Papaspiliopoulos and Giacomo Zanella. 2025. "Partially Factorized Variational Inference for High-Dimensional Mixed Models." *Biometrika* 112(2):asae067.
- Grossmann, Matt, Marty P. Jordan and Joshua McCrain. 2021. "The Correlates of State Policy and the Structure of State Panel Data." *State Politics & Policy Quarterly* 21(4):430–450.
- Heeringa, Steven G., Judith H. Connor and Doris C. Darrah. 1986. 1980 SRC National Sample: Design and Development. Technical report Institute for Social Research, The University of Michigan Ann Arbor, Mich.: .
- Hopkins, Daniel J. 2018. *The Increasingly United States: How and Why American Political Behavior Nationalized*. Chicago: University of Chicago Press.
- Jacobson, Gary C. 2015. "It's Nothing Personal: The Decline of the Incumbency Advantage in US House Elections." *The Journal of Politics* 77(3):861–873.
- Johnson, Martin, Paul Brace and Kevin Arceneaux. 2005. "Public Opinion and Dynamic

- Representation in the American States: The Case of Environmental Attitudes.” *Social Science Quarterly* 86(1):87–108.
- Kish, Leslie. 1965. *Survey Sampling*. New York: Wiley.
- Kish, Leslie and Irene Hess. 1965. The Survey Research Center’s National Sample of Dwellings. Technical Report 2315 Institute for Social Research, The University of Michigan Ann Arbor, Mich.: .
- Kuriwaki, Shiro. 2020. “Ticket Splitting in a Nationalized Era.” Working paper.
- Lagodny, Julius, Rebekah Jones, Julianna Koch and Peter K. Enns. 2023. “A Validation and Extension of State-Level Public Policy Mood: 1956–2020.” *State Politics & Policy Quarterly* 23(4):359–372.
- Lax, Jeffrey R. and Justin H. Phillips. 2009a. “Gay Rights in the States: Public Opinion and Policy Responsiveness.” *American Political Science Review* 103(3):367–386.
- Lax, Jeffrey R. and Justin H. Phillips. 2009b. “How Should We Estimate Public Opinion in The States?” *American Journal of Political Science* 53(1):107–121.
- Lax, Jeffrey R. and Justin H. Phillips. 2012. “The Democratic Deficit in the States.” *American Journal of Political Science* 56(1):148–166.
- Lax, Jeffrey R., Justin H. Phillips and Adam Zelizer. 2019. “The Party or the Purse? Unequal Representation in the US Senate.” *American Political Science Review* 113(4):917–940.
- Matsusaka, John G. 2001. “Problems with a Methodology Used to Evaluate the Voter Initiative.” *The Journal of Politics* 63(4):1250–1256.
- Mehlhoff, Isaac D. 2024. “Leveraging Time-Series Information to Improve Small-Area Estimation.” Working paper.
- Norrander, Barbara. 2001. “Measuring State Public Opinion with the Senate National Election Study.” *State Politics & Policy Quarterly* 1(1):111–125.
- Ornstein, Joseph T. 2020. “Stacked Regression and Poststratification.” *Political Analysis* 28(2):293–301.

- Pacheco, Julianna. 2014. "Measuring and Evaluating Changes in State Opinion Across Eight Issues." *American Politics Research* 42(6):986–1009.
- Page, Benjamin I. and Robert Y. Shapiro. 1992. *The Rational Public: Fifty Years of Trends in Americans' Policy Preferences*. American Politics and Political Economy Series Chicago: University of Chicago Press.
- Park, David K., Andrew Gelman and Joseph Bafumi. 2004. "Bayesian Multilevel Estimation with Poststratification: State-Level Estimates from National Polls." *Political Analysis* 12(4):375–385.
- Rogers, Steven. 2016. "National Forces in State Legislative Elections." *The ANNALS of the American Academy of Political and Social Science* 667(1):207–225.
- Shirley, Kenneth E. and Andrew Gelman. 2015. "Hierarchical Models for Estimating State and Demographic Trends in US Death Penalty Public Opinion." *Journal of the Royal Statistical Society. Series A (Statistics in Society)* 178(1):1–28.
- Sievert, Joel and Seth C. McKee. 2019. "Nationalization in U.S. Senate and Gubernatorial Elections." *American Politics Research* 47(5):1055–1080.
- Stimson, James A., Michael B. Mackuen and Robert S. Erikson. 1995. "Dynamic Representation." *American Political Science Review* 89(3):543–565.
- Stollwerk, Alissa F. 2017. *Essays on the Measurement of Public Opinion* PhD thesis Columbia University.
- Tausanovitch, Chris and Christopher Warshaw. 2013. "Measuring Constituent Policy Preferences in Congress, State Legislatures, and Cities." *The Journal of Politics* 75(2):330–342.
- Warshaw, Christopher and Jonathan Rodden. 2012. "How Should We Measure District-Level Public Opinion on Individual Issues?" *The Journal of Politics* 74(1):203–219.
- Wright, Gerald C. 1989. "Policy Voting in the U. S. Senate: Who Is Represented?" *Legislative Studies Quarterly* 14(4):465–486.