

Public Opinion in Subnational Politics

Devin Caughey, MIT

Christopher Warshaw, George Washington University

Until recently, the study of representation at the subnational level was hobbled by the lack of high-quality information about public opinion. The advent of new data sources, however, as well as of new methods such as multilevel regression and poststratification, has greatly enhanced scholars' capacity to describe public opinion in states, legislative districts, cities, and other subnational units. These advances in measurement have in turn revolutionized the study of subnational representation. In this article, we summarize new approaches to the measurement of subnational opinion. We then review recent developments in the study of the role of subnational public opinion in the political process and discuss potentially fruitful avenues for future research.

The advent of new data sources and statistical techniques has led to a revolution in scholars' ability to examine public opinion and representation at the subnational level. The most important methodological change has been the development of regularization methods for subnational opinion estimation, particularly multilevel regression and poststratification, or MRP (Gelman and Little 1997). Thanks to these developments, the study of public opinion and representation in subnational politics has rapidly advanced in recent years. Multiple studies on public opinion and representation in American state governments and cities have appeared in each of the discipline's flagship journals over the past decade. In this article, we provide an overview of how scholars can measure public opinion at the subnational level. We illustrate the discussion through an application to public opinion and representation on gay rights. We discuss recent substantive advances in the study of representation at the subnational level. Finally, we discuss where we see the study of public opinion in state and local politics going in the coming years.

NEW DATA ON PUBLIC OPINION

One of the most important reasons that the study of public opinion and representation in subnational politics has grown dramatically over the past decade has been the availability

of new data sources. Much of this progress has been due to the efforts of the Roper Center for Public Opinion (<https://ropercenter.cornell.edu>). The center has collected over 20,000 polls from hundreds of survey firms going back to the 1930s and has made the individual-level data available for download by researchers from member institutions. These commercial polls offer an enormous wealth of data on a huge variety of topics. For example, the archive contains over 100 polls from the past two decades that cover the topic of gay rights, our running example in this article.

In addition, a number of large-scale academic surveys have been developed over the past two decades that examine public opinion on dozens of individual issues. In 2000, and then again 2004 and 2008, researchers at the University of Pennsylvania's Annenberg School of Communications surveyed over 50,000 people over the course of the presidential election campaign. In recent years, cooperative surveys have emerged that produce large national samples by aggregating numerous small sample surveys, most notably the Cooperative Congressional Election Study and the Cooperative Campaign Analysis Project. The combination of all these data sources means that scholars can now examine the causes and consequences of public opinion on hundreds of issues over the past three-quarters of a century.

Devin Caughey (devin.caughey@gmail.com) is the Silverman (1968) Family Career Development Associate Professor (without tenure) in the Department of Political Science, Massachusetts Institute of Technology, Cambridge, MA 02139. Christopher Warshaw (cwarshaw@gmail.com) is an assistant professor in the Department of Political Science, George Washington University, Washington, DC 20052.

This article was solicited by the editors, was presented and critiqued in a workshop, and went through editorial review. Data and supporting materials necessary to reproduce the numerical results in the article are available in the *JOP* Dataverse (<https://dataverse.harvard.edu/dataverse/jop>).

ESTIMATING SUBNATIONAL OPINION

Of course, new data sources have not been enough on their own. Rather, it has been the combination of new data and methods that has really kick-started work on subnational opinion and representation. Through the end of the twentieth century, political scientists' primary approach to measuring subnational public opinion was to "disaggregate" one or more national surveys and take the average (possibly accounting for sampling weights) in each subnational unit (e.g., Brace et al. 2002; Erikson, Wright, and McIver 1993; Miller and Stokes 1963). Since the publication of Gelman and Little's (1997) seminal MRP article, however, political scientists have increasingly turned to methods that combine model-based regularization with post hoc weighting.

MRP and its relatives have been shown to perform well on samples as small as a few thousand people and to have lower cross-validated prediction error than disaggregation (Lax and Phillips 2009b; Park, Gelman, and Bafumi 2004; Warshaw and Rodden 2012; but see Buttice and Highton 2013). Indeed, some have gone so far as to call MRP the "gold standard" for estimating subnational opinion (Selb and Munzert 2011, 456). Our position is more nuanced. In some contexts, especially when public opinion is the dependent variable and unbiasedness is more important than efficiency, design-based disaggregation can be preferable to model-based methods such as MRP. In many others, the advantages of a model-based approach outweigh its drawbacks, but even then the form of regularization and other modeling choices must be carefully considered. In this section, we discuss various approaches to subnational opinion estimation and illustrate trade-offs between them using the well-studied example of public opinion on gay rights.

Disaggregation and other design-based methods

The most straightforward estimator of subnational opinion is the sample average in a given subnational unit. The sample average is a "design-based" estimator in that its validity can be grounded solely in assumptions about the sampling design (e.g., that the data are a simple random sample) rather than in a model of the data-generating process in the population (e.g., that opinion is normally distributed, conditional on a linear combination of covariates). By incorporating sampling weights, unbiased design-based estimators can be derived for almost any sampling design, including those in which sampling probabilities are correlated or unequal across units (see, e.g., Lumley 2010).

Design-based estimators have the advantage of being transparent and easy to calculate and of not relying on difficult-to-verify parametric assumptions. It is thus unsurprising that almost every study of subnational opinion conducted before

the 2000s relied on design-based estimators, typically using data disaggregated from national surveys. Design-based estimators have important drawbacks, however. First, because of clustered sampling designs, nonresponse bias, and other factors, subnational survey samples are often unrepresentative, thus biasing nominally unbiased estimators. Second, disaggregating from national surveys often results in very small samples at the subnational level, yielding highly imprecise estimates. In principle, the first problem can be ameliorated with poststratification or other weighting methods (Särndal and Lundström 2005), and the second, by pooling together surveys to increase sample sizes (Erikson et al. 1993). Ultimately, however, the goals of reducing bias and increasing precision are in tension with one another. Increasing the number of poststratification variables, for example, often increases the variance of estimators (Little and Vartivarian 2005) and, in the limit, leads inevitably to some population cells being absent from the sample. As we discuss later, MRP can be thought of as a model-based method for managing this trade-off between bias and variance.

Despite these limitations, the usefulness of disaggregation should not be overlooked. When subnational samples are close to random and not too small, a design-based estimator, possibly combined with adjustment weighting, should be approximately unbiased and reasonably precise. Further, when analyzing opinion as an outcome variable, researchers might well prefer a somewhat noisy but unbiased design-based measure than a model-based one that trades off bias for lower variance.¹ In particular, we suggest that studies of causal effects on public opinion (e.g., policy feedback) should generally use disaggregation rather than a model-based method such as MRP.

To illustrate the trade-offs between methods, we examine the policy feedback effects of the Supreme Court's decision in *Obergefell v. Hodges* to legalize same-sex marriage in 2015. There is evidence that exposure to gay people has driven changes in public opinion on gay rights. It thus stands to reason that observing same-sex marriages, or married gay couples, might lead people to change their views on same-sex marriage (Movement Advancement Project 2018). On the eve of the Supreme Court decision, 13 states did not allow same-sex marriage gay rights. Did the public in those states change their views on same-sex marriage in the wake of the Supreme Court decision? To examine this question, we use a simple difference-in-differences model that examines changes in state-level public opinion between 2015 and 2016 (table 1).

1. Measurement error in a dependent variable inflates standard errors but does not usually bias regression estimates (Lewis and Linzer 2005).

Table 1. Policy Feedback Effect of Legalization of Same-Sex Marriage on Opinion

| | Disaggregation Large Sample (Unweighted) (1) | Disaggregation Small Sample (Unweighted) (2) | Disaggregation Small Sample (Weighted) (3) | Multilevel Regression and Poststratification Small Sample (Weighted) (4) |
|----------------------------|---|---|---|--|
| Same-sex marriage legal | .034** (.013) | .030 (.017) | .018 (.018) | .006*** (.001) |
| State fixed effects | Yes | Yes | Yes | Yes |
| Year fixed effects | Yes | Yes | Yes | Yes |
| R ² | .951 | .942 | .936 | 1.000 |

Note. Standard errors (in parentheses) are clustered by state. $N = 100$.

** $p < .01$.

*** $p < .001$.

In order to evaluate how data and modeling decisions affect substantive findings on policy feedback effects, we compare several measurement choices. First, we use a massive set of survey data from the Public Religion Research Institute's American Values Atlas. In 2015 and 2016, they interviewed approximately 50,000 Americans in each year. This provides large-scale samples in each state and enables us to generate relatively precise estimates of public opinion on same-sex marriage in each state (although it is important to note that their samples are unweighted within each state). In column 1 of table 1, we find that in states where same-sex marriage was legalized, public opinion shifted about 3 points in favor of same-sex marriage. Moreover, this effect is significant at the .01 level.

Next, we consider a smaller set of public opinion data using surveys that we downloaded from the Roper Center. This provided a sample of approximately 9,000 Americans in 2015 and 10,000 in 2016. Column 2 of table 1 shows the results using unweighted disaggregated estimates of opinion in each state. It indicates that the substantive size of the effect of the Supreme Court decision on public opinion is very similar to the effect in column 1. But the effect is estimated much more noisily. Column 3 shows a similar result using weighted disaggregated estimates.

Finally, we used a dynamic MRP model to smooth public opinion across states.² This model increased the accuracy of the estimates of public opinion in each state. But it also smoothed away most of the treatment effect. Indeed, column 4 of table 1 shows that MRP effect of the Supreme Court decision is estimated to be significantly smaller than the large-

sample disaggregation estimate in column 1 (although still distinguishable from 0). Note also that the model now explains essentially all the variation in the smoothed estimates.

This example illustrates the bias-variance trade-off between disaggregation and smoothing models. Under the assumption of simple random sampling, the disaggregated samples in columns 1–3 yield unbiased estimates of the policy feedback effects of gay marriage policies. However, the point estimates from these models are imprecise—especially in the models in columns 2 and 3 that draw on small disaggregated opinion samples from the handful of surveys available in those years via the Roper Center. In contrast, the estimate of policy feedback effects in column 4 that uses an MRP model is very precisely estimated. But the policy feedback effect has been attenuated to close to zero.

MRP and other model-based regularization methods

As noted above, MRP was developed as a means of addressing the twin problems of bias and variance, which respectively derive from the unrepresentativeness and small size of many subnational survey samples. MRP entails two steps. First, a multilevel regression model is used to estimate opinion in population cells defined by the cross-classification of geographic and demographic variables (e.g., state, race, and gender). Second, opinion in each subnational unit is estimated by poststratifying (i.e., weighting) the cell estimates in proportion to their share of the subnational population. Because the multilevel model regularizes each cell estimate by “shrinking” its estimate toward observably similar cells, the model increases the estimates' precision at the expense of some increase in bias.

2. This model is discussed in more detail below.

Of course, MRP is only one possible way of combining regularized prediction and post hoc weighting. An alternative regularization method, such as the lasso or Bayesian additive regression trees, could be used (e.g., Bisbee 2018; Caughey and Hartman 2017; Goplerud et al. 2018). Even within the multilevel regression framework, one must choose from which observations to “borrow strength.” In classic MRP, the smoothing is purely cross-sectional. Two other options are to borrow strength over time through a dynamic model or from responses to other survey questions through a latent-variable model. Since straightforward poststratification is not always feasible, an alternative approach to weighting the cell estimates may be used. Different approaches may be optimal for different purposes, and each of these choices is a subject of ongoing research.

Smoothing cross-sectionally. The first-generation MRP models developed by Park et al. (2004) borrow strength cross-sectionally. That is, the estimate for a given population cell is smoothed using data from similar states or demographic groups. For example, the opinion estimate for black women in Mississippi might be informed by the responses of other Mississippians as well as by those of women and African Americans in other states. Moreover, through a second-level model, the estimate for the average Mississippian might be influenced by the estimates from similar states, such as Alabama.

To illustrate, we use the MRP model developed by Lax and Phillips (2009a, 2009b). This model uses a multilevel

logistic regression to smooth state-level opinion on gay marriage and other gay rights issues cross-sectionally for a set of demographic-geographic strata. The level 1 model for the response of individual i is

$$\Pr(y_i = 1) = \text{logit}^{-1}(\gamma_0 + \alpha_{r[i]}^{\text{race}} + \alpha_{g[i]}^{\text{gender}} + \alpha_{e[i]}^{\text{edu}} + \alpha_{a[i]}^{\text{age}} + \alpha_{s[i]}^{\text{state}} + \alpha_{p[i]}^{\text{poll}}), \quad (1)$$

for $r \in \{1, \dots, 4\}$, $g \in \{1, 2\}$, $e \in \{1, \dots, 5\}$, $a \in \{1, \dots, 4\}$, and $p \in \{1, \dots, \text{no. polls}\}$. In the second level of the model, the state effects α^{state} are modeled as a function of the region into which the state falls, the state’s percentage of evangelical or Mormon residents, and the Democratic presidential vote share in the last election:

$$\alpha_s^{\text{state}} \sim N(\alpha_z^{\text{region}} + \beta_1 \times \text{pres}_p + \beta_2 \times \text{religion}_s, \sigma_s^2), \quad (2)$$

where $s \in \{1, \dots, 51\}$

We estimate a slightly simpler version of this model using the dgo package in R (Dunham, Caughey, and Warshaw 2016). In our model, we estimate public opinion in each state on same-sex marriage in 2012 using age and education as individual-level strata and religion, union membership, median income, and the percentage of each state’s residents in same-sex relationships as state-level predictors. The results are shown in figure 1. In addition, in figure 2 we replicate Lax and Phillips’s (2009a) finding of a strong relationship between mass support for same-sex marriage in 2012 and whether a state sanctioned same-sex marriage in 2013.




Figure 1. Mass support for same-sex marriage (2012). Lighter shades denote higher support




Figure 2. Policy responsiveness to opinion on same-sex marriage (2012–13)

We can do a similar analysis to show that public opinion on same-sex marriage is associated with municipal policies on same-sex marriage (Warshaw 2016). Figure 3 shows that there is a strong correlation between the rights that municipalities grant to gay employees and public opinion on same-sex marriage.³ Cities with greater support for same-sex marriage are much more likely to provide strong protections to gay employees. This suggests that city governments are responsive to the views of their citizens on gay rights (see also Einstein and Kogan 2016; Palus 2010; Tausanovitch and Warshaw 2014).

There are a number of best practices that scholars have developed for cross-sectional MRP models. First, MRP often performs well with samples as small as a few thousand survey respondents (Lax and Phillips 2009b; Warshaw and Rodden 2012), but its performance is heterogeneous across issues (Buttice and Highton 2013). MRP is almost always more accurate as sample sizes increase. Second, scholars should generally use at least one variable to help predict opinion at each geographic level in the multilevel model (Lax and Phillips 2013). Moreover, researchers should spend time making sure that these variables are good predictors of the variation in opinion across geography (Buttice and Highton 2013). Most of the performance gains in smoothing models come from the inclusion of good constituency-level predictors (Hanretty, Lauderdale, and Vivyan 2018; Warshaw and Rodden 2012). Finally, smoothing based on geographically proximate units

can also improve predictive accuracy and compensate for weak constituency-level predictors (Hanretty et al. 2018).

Future work is likely to use machine learning approaches to smooth public opinion based on more than the handful of predictors leveraged in existing work (e.g., Bisbee 2018; Caughey and Hartman 2017; Goplerud et al. 2018). This work shows great promise for further improving the ability of scholars to develop accurate estimates of the cross-sectional variation in public opinion at the subnational level.

A more foundational limitation of cross-sectional public opinion models is that they are ill suited for determining the causal effect of public opinion on public policy and other important outcomes (Brace 2018; Lowery, Gray, and Hager 1989). It is impossible to rule out reverse causation whereby public opinion is influenced by a policy feedback process. For instance, figures 2 and 3 show that public opinion and policy on same-sex marriage are correlated. But they cannot rule out the possibility that the order of causality is reversed, and the establishment of same-sex marriage leads to an increase in support for gay marriage in the mass public. In addition, some omitted variable could be confounding the relationship between public opinion and policy. For example, perhaps partisan turnover (e.g., the election of Democratic governors) was the key factor that led to the legalization of same-sex marriage. One way to bolster causal inferences is to switch from a cross-sectional to a dynamic perspective and examine differences in public opinion across time.

Smoothing over time. Examining subnational opinion over time raises new questions about whether to borrow strength cross-sectionally from other units at the same point in time or instead (or in addition) to do so dynamically, from data on the same unit at different points in time. At one end of the spectrum is the option of estimating separate cross-sectional

3. The data on the rights that municipalities grant to gay employees are derived from the Human Rights Campaign's Municipal Equality Index, which scores how well approximately 300 cities support the LGBT people who live and work there.




Figure 3. City responsiveness to public opinion on gay rights: association between public opinion on same-sex marriage in each city and an index of the rights that municipalities grant to gay employees.

MRP models in every year (e.g., Enns and Koch 2013; Lewis and Jacobsmeier 2017). The advantage of this approach is that it maximizes flexibility and variation over time. Its main disadvantage is the other side of the coin: years when data are thin will have imprecise estimates, and when data in a particular time period are missing altogether estimates cannot be produced.

The other extreme, smoothing across time only, is predicated on the often reasonable assumption that the best guess for opinion in one year is opinion on the same question in the years immediately before or after. We would expect Democratic party identification (PID) in New Hampshire in 1957, for instance, to be strongly predicted by its PID in 1956—more

strongly, perhaps, than by 1957 PID in similar states such as Maine or Vermont. Such substantive information can be encoded in various ways, such as through a moving average (Pacheco 2011), a linear or quadratic time trend (Gelman et al. 2016; Shirley and Gelman 2015), or a Bayesian dynamic linear model (DLM; Caughey and Warshaw 2015, 2018; Linzer 2013).

In between these two extremes, models that smooth over time can be combined with cross-sectional information. Pacheco (2011), for example, first applies cross-sectional MRP in each year separately and then takes a moving average of the estimates over time. A disadvantage of this two-step approach is that it does not automatically propagate uncertainty in the first-stage estimates through to the second

stage. An alternative is to estimate a unified Bayesian model in which the cross-sectional parameters are smoothed over time (Caughey and Warshaw 2015) or in which opinion is modeled with a DLM that incorporates both lagged opinion and cross-sectional covariates (Caughey and Warshaw 2018). Estimating a unified model properly accounts for uncertainty, although at the possible expense of tractability and computation time.

We use this last approach to estimate public opinion on same-sex marriage in each state-year from 1999 to 2016. Figure 4 shows that public support for same-sex marriage has increased in every state during this time period, although with some heterogeneity in the time trends across states. To examine whether these state-specific trends were associated with the adoption of same-sex marriage laws, we estimate a simple regression model in which the outcome takes a value of 0 in states that ban same-sex marriage and civil unions, 1

in states that allow civil unions, and 2 in states that allow same sex marriage.

Table 2 compares the results of this analysis with the same analysis using disaggregated estimates for each year. First, columns 1 and 2 show the cross-sectional association between mass opinion on same-sex marriage and policy (averaged across the 1999–2015 period). Consistent with the findings in Lax and Phillips (2009a), both the disaggregated (col. 1) and smoothed (col. 2) estimates of public opinion on same-sex marriage have a strong cross-sectional association with policy. But the greater measurement error in the disaggregated estimates attenuates their relationship with policy compared to the smoothed estimates.

Next, in table 2 columns 3 and 4, we estimate models with fixed effects for state and year to account for state- and time-invariant confounders, thus leveraging only within-state variation over time. Because of the measurement error in the disaggregated estimates, the regression estimate of their dynamic effect on policy is attenuated to essentially 0 (col. 3). However, the smoothed estimates of opinion have a large and significant effect on public policy (col. 4). This evidence is reassuring about the performance of statehouse democracy on same-sex marriage and reinforces the dynamic findings about responsiveness on same-sex marriage in Lewis and Jacobsmeier (2017).

Modeling public opinion as a latent variable. Latent-variable models (LVMs), which model multiple observed variables as functions of one or more unobserved ones, can be thought of as an especially elaborate form of model-based prediction. The advantage of an LVM is that it combines information not only from multiple units or time points but also from multiple observed variables. In an item-response theory (IRT) model, for example, the probability of selecting the liberal response option to a given dichotomous question (e.g., favoring same-sex marriage) is modeled as function of a latent variable (e.g., gay-rights liberalism), the question’s “discrimination” with respect to that latent variable, and a “difficulty” parameter that reflects baseline support for the question.

LVMs can also be combined with other forms of smoothing. Tausanovitch and Warshaw (2013), for example, use an IRT model to estimate the liberalism of thousands of geocoded survey respondents, then apply cross-sectional MRP to the liberalism estimates to produce a measure of average liberalism in states, districts, and cities. Similarly, Caughey and Warshaw (2018) smooth IRT estimates of latent liberalism across time using a DLM.

The target of inference in LVMs is typically the latent variable itself, rather than any one of its indicators. Predictions for individual indicators, however, can easily be gen-




Figure 4. Change in public support for same-sex marriage between 1999 (left) and 2015 (right).

Table 2. Dynamic Responsiveness on Same-Sex Marriage

| | Disaggregated (CS) (1) | Smoothed (CS) (2) | Disaggregated (TSCS) (3) | Smoothed (TSCS) (4) |
|-----------------------------------|---------------------------|----------------------|-----------------------------|------------------------|
| Support for same-sex marriage (%) | .013*** (.003) | .027*** (.006) | .001 (.001) | .113*** (.025) |
| R^2 | .210 | .376 | .610 | .672 |
| State fixed effects | | | Yes | Yes |
| Year fixed effects | Yes | Yes | Yes | Yes |

Note. Standard errors (in parentheses) are clustered by state. CS = cross-sectional; TS = time and state. $N = 584$.

*** $p < .001$.

erated using the estimated values of the latent variable and the parameters of the model. Alternatively, latent-variable scores can be included in a multilevel model predicting individual question responses (e.g., Kestellec 2018). In either case, the estimates for specific survey questions could well be improved by the information from related questions (e.g., for same-sex marriage, questions on antidiscrimination protections for homosexuals). The LVM could even be used to impute the margins on a specific question for years or subnational units in which that specific question was not asked. Although as yet rare, the use of LVMs to improve subnational estimates for specific questions is a promising area for future research.

Weighting. The canonical smoothing (MRP) models developed by Park et al. (2004) use simple poststratification weights to sum the predictions for each geographic-demographic strata and develop estimates of public opinion in each geographic area. These models typically assume that the population targets used to generate poststratification weights for each geographic demographic are known with certainty, usually based on census data. Moreover, they implicitly assume that respondents are sampled randomly within demographic strata. Of course, these two requirements are often unfulfilled in real-life applications.

In regard to the first issue, the canonical MRP models require scholars to have poststratification weights for each geographic-demographic strata. In the United States, it is often straightforward to obtain this information at the state level. But it is difficult to obtain these types of detailed population breakdowns below the state level. And in other countries, this kind of detailed census microdata is often completely unavailable (Leemann and Wasserfallen 2017). The only available census data are often just the marginals of different groups in the state population (e.g., the percentage of people that are women). This has made it difficult for

comparative politics scholars to deploy smoothing models such as MRP to study representation and opinion.

There are a number of possible approaches to address this problem. In situations in which population margins are available, Leemann and Wasserfallen (2017) show that it is possible to generate poststratification weights based on aggregate-level census data using an approach similar to raking. In situations in which population margins are not even known with certainty, such as the distribution of partisans across states, scholars could use more complicated modeling approaches to generate population targets that combine different sources of available information (Caughey and Wang 2018; Kestellec et al. 2015; Kiewiet de Jonge, Langer, and Sinozich 2018).

Recent work has also begun to address several other lingering issues related to weighting the estimates of a smoothing model. First, sampling in surveys is often not random within strata, such as in cluster-sampled surveys.⁴ In many cases, this requires analysts to either generate their own within-strata data weights or use weights provided by the survey firm (see Caughey and Warshaw 2015; Ghitza and Gelman 2013). There is also a question about how many poststratification cells to use. A tentative lesson of recent work is that scholars usually do not need to use a complicated first-stage model with a large number of demographic strata (groups) when the substantive goal is to estimate public opinion at the geo-

4. Many surveys use area-sampling designs in which respondents are clustered within geographic areas. In this approach, researchers sample a set of primary sampling units (PSUs), such as cities or census tracts. Then they interview a sample of people within each PSU. This approach yields samples that may be unbiased at the national level. But they could produce a nonrepresentative selection of PSUs for any particular geographic subunit. For instance, the PSU for Wyoming could be centered around Jackson Hole. Cluster sampling was standard practice in commercial polls before the 1970s, and it continues to be used by many high-quality academic surveys, including the General Social Survey and the American National Election Survey.

graphic level (Lax and Phillips 2013). However, more complicated first-stage models can be useful when the goal is to model the opinion of small demographic groups within states or other geographic units (Ghitza and Gelman 2013).

Uncertainty in public opinion estimates

A final important issue in the measurement and usage of public opinion at the subnational level is that scholars need to be mindful that there is always uncertainty in public opinion estimates.⁵ It is important for scholars to take this uncertainty into account in any substantive analysis (Achen 1978; Kestellec et al. 2015; Lax and Phillips 2013). It is particularly crucial when the estimates of public opinion are used as a predictor in a regression model (see Gelman and Hill 2006, 542). In addition to traditional errors-in-variables corrections, two helpful frameworks for propagating measurement error into substantive inferences are the “method of composition” (Treier and Jackman 2008, 215–16), which entails drawing samples from the joint posterior distribution of the measurement and analysis models, and “multiple overimputation” (Blackwell, Honaker, and King 2017), which applies a correction to the sampling variance analogous to that used for multiple imputation. Several recent papers have used such approaches to propagate the uncertainty in their estimates of public opinion into their substantive models (Caughey and Warshaw 2018; Kestellec et al. 2015).⁶

PUBLIC OPINION AND REPRESENTATION

The geography of subnational public opinion

One of the most important tasks of public opinion research is to describe geographic variation in the mass public’s views. There is a large literature that has used recent methodological advances in our ability to measure subnational public opinion to examine geographic variation in public opinion. In many cases, this work forms the foundation for substantive projects. For instance, Lax and Phillips (2009a) use MRP to measure state-level public opinion on gay-rights issues in the 2000s and examine the association between public opinion and policy. In other cases, however, public opinion is presented for purely descriptive purposes. Tausanovitch and Warshaw (2013) develop estimates of the policy ideology of every city and legislative district in the country. Elmendorf and Spencer (2014) estimate the average level of racial prejudice in every state and county in the country. They find the highest levels of racial prejudice in southern states such as

Mississippi and South Carolina. However, they also find high levels of racial prejudice in several other states, such as Wyoming, Pennsylvania, and Ohio. Their findings provide policy makers with information about contemporary levels of racial prejudice in the United States that could be useful for future revisions to the Voting Rights Act and other laws protecting minorities.

The growing availability of thousands of surveys over the entire span of the last half century has revolutionized scholars’ ability to study variation in public opinion in earlier time periods. Caughey (2018), for example, combines data from hundreds of polls from 1936 to 1952 to estimate support for New Deal liberalism in each state (see also Krimmel and Rader 2016). More generally, scholars have taken advantage of the lengthening time span of available survey data to examine trends in subnational opinion over many years. Examples include Pacheco’s (2014) 1977–2004 study of state-level support for spending in various policy areas, Shirley and Gelman’s (2015) 50-year study of state opinion on gun control, and Caughey and Warshaw’s (2018) and Enns and Koch’s (2013) estimates of ideological trends in the states over the past 75 years.

Policy representation

The recent advances in our ability to model public opinion have led to three major advances in the study of representation. First, it has enabled scholars to study representation at new geographic levels. Table 3 shows a sample of the recent studies that have used smoothing models to examine representation at various geographic levels in both the United States and other countries. A number of studies have examined representation at the state level, while a smaller number have examined representation below the state level. We expect much work in the coming years to focus on local governments in the United States (e.g., school districts) and increasingly to focus on comparative politics outside the United States.⁷ We also expect more work to focus on dyadic representation in legislatures. Second, the recent advances in opinion estimation have enabled scholars to study representation on many individual issues. For instance, Lax and Phillips (2012) estimate state-level public opinion on dozens of individual issues in the 2000s and examine whether public policy on these issues is responsive and congruent with pub-

5. This is true whether the estimates are generated using disaggregation, MRP, or some other model-based approach.

6. In our experience, correcting for measurement error often attenuates the estimated effects of public opinion relative to the unadjusted estimates.

7. However, data limitations are likely to continue to complicate the ability of scholars to study representation in new contexts. For example, substate geographic identifiers are often unavailable in survey data in the United States before about the year 2000. Moreover, subnational geographic identifiers are only sparsely available in surveys outside the United States.

Table 3. Selection of Studies of Representation at Subnational Level

| Study | Level of Public Opinion | Topic | XS/Dynamic |
|--------------------------------------|--------------------------------------|----------------------------|------------|
| Tausanovitch and Warshaw (2014) | American cities | Policy representation | XS |
| Berkman and Plutzer (2005) | American schools | Policy representation | XS |
| Lax and Phillips (2009b) | American states | Policy representation | XS |
| Lax and Phillips (2012) | American states | Policy representation | XS |
| Pacheco (2013) | American states | Policy representation | Dynamic |
| Caughey and Warshaw (2018) | American states | Policy representation | Dynamic |
| Pacheco (2012) | American states | Policy representation | XS |
| Lewis and Jacobsmeier (2017) | American states | Policy representation | Dynamic |
| Canes-Wrone, Clark, and Kelly (2014) | American states | Judicial representation | XS |
| Kastellec, Lax, and Phillips (2010) | American states | Legislative representation | XS |
| Kastellec et al. (2015) | American state parties | Legislative representation | XS |
| Shor (2017) | American state legislative districts | Legislative representation | XS |
| Broockman and Skovron (2018) | American state legislative districts | Elite perceptions | XS |
| Leemann and Wasserfallen (2016) | Swiss states | Policy representation | XS |

Note. XS = cross-sectional.

lic opinion. Overall, they find a strong link between public opinion and policy. Finally, the availability of models that smooth opinion over time has enabled scholars to examine the dynamic effects of public opinion on public policy and other important outcomes. Given the move toward causal inference in the discipline, we expect that the study of dynamic responsiveness is going to increasingly dominate the study of representation.

Policy feedback

There is a smaller literature that has examined policy feedback effects on subnational public opinion. For example, Sances and Clinton (2017) examine whether the Medicaid expansion in the Affordable Care Act (ACA) affected public opinion. They find that the ACA had only small effects on opinion. But it is hard to know whether Sances and Clinton's (2017) findings represent the modal effect of policy on public opinion or whether the effects of the ACA are unusually small. The growing availability of historical opinion data and the recent development of large-scale cooperative surveys enable scholars to study policy feedback effects in a variety of new contexts.

CONCLUSION

Until recently, the lack of high-quality information about public opinion at the subnational level was an important barrier that hindered the study of representation in Congress, state government, and local government. In recent years, the availability of new data sources and the advent of smoothing models to characterize mass opinion at the subnational level has revolutionized scholars' ability to describe public opinion and

examine its influence on the political process in a variety of important contexts.

ACKNOWLEDGMENTS

We thank Chris Tausanovitch and participants in the 2018 PIPE Symposium on Studying Subnational Policy Making for feedback. All mistakes, however, are our own.

REFERENCES

- Achen, Christopher H. 1978. "Measuring Representation." *American Journal of Political Science* 22 (3): 475–510.
- Berkman, Michael B., and Eric Plutzer. 2005. *Ten Thousand Democracies: Politics and Public Opinion in America's School Districts*. Washington, DC: Georgetown University Press.
- Bisbee, James. 2018. "BARP: Improving Mister P using Bayesian Additive Regression Trees." Unpublished manuscript.
- Blackwell, Matthew, James Honaker, and Gary King. 2017. "A Unified Approach to Measurement Error and Missing Data: Overview and Applications." *Sociological Methods and Research* 46 (3): 303–41.
- Brace, Paul. 2018. "Aggregating Survey Data to Estimate Subnational Public Opinion." In Lonna Rae Atkeson and R. Michael Alvarez, eds., *The Oxford Handbook of Polling and Survey Methods*. New York: Oxford University Press.
- 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* 46:173–89.
- Broockman, David E., and Christopher Skovron. 2018. "Bias in Perceptions of Public Opinion among Political Elites." *American Political Science Review* 112 (3): 542–63.
- Buttice, Matthew K., and Benjamin Highton. 2013. "How Does Multilevel Regression and Poststratification Perform with Conventional National Surveys?" *Political Analysis* 21 (4): 449–67.
- Canes-Wrone, Brandice, Tom S. Clark, and Jason P. Kelly. 2014. "Judicial Selection and Death Penalty Decisions." *American Political Science Review* 108 (1): 23–39.

- Caughey, Devin. 2018. *The Unsolid South: Mass Politics and National Representation in a One-Party Enclave*. Princeton, NJ: Princeton University Press.
- Caughey, Devin, and Erin Hartman. 2017. "Target Selection as Variable Selection: Using the Lasso to Select Auxiliary Vectors for the Construction of Survey Weights." Unpublished manuscript.
- Caughey, Devin, and Mallory Wang. 2018. "Dynamic Ecological Inference for Time-Varying Population Distributions Based on Sparse, Irregular, and Noisy Marginal Data." *Political Analysis* (forthcoming).
- 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): 249–66.
- Dunham, James, Devin Caughey, and Christopher Warshaw. 2016. dgo: Dynamic Estimation of Group-Level Opinion. R package version 0.2.3. <https://jamesdunham.github.io/dgo/>.
- Einstein, Katherine Levine, and Vladimir Kogan. 2016. "Pushing the City Limits: Policy Responsiveness in Municipal Government." *Urban Affairs Review* 52 (1): 3–32.
- Elmendorf, Christopher S., and Douglas M. Spencer. 2014. "The Geography of Racial Stereotyping: Evidence and Implications for VRA Preclearance after Shelby County." *California Law Review* 102:1123–80.
- Enns, Peter K., and Julianna Koch. 2013. "Public Opinion in the U.S. States, 1956 to 2010." *State Politics and Policy Quarterly* 13 (3): 349–72.
- 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.
- Gelman, Andrew, and Jennifer Hill. 2006. *Data Analysis Using Regression and Hierarchical/Multilevel Models*. New York: Cambridge University Press.
- Gelman, Andrew, Jeffrey Lax, Justin Phillips, Jonah Gabry, and Robert Trangucci. 2016. "Using Multilevel Regression and Poststratification to Estimate Dynamic Public Opinion." Unpublished manuscript, Columbia University.
- 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): 762–76.
- Goplerud, Max, Shiro Kuriwaki, Marc Ratkovic, and Dustin Tingley. 2018. "Sparse Multilevel Regression and Poststratification (sMRP)." Unpublished manuscript, Harvard University.
- Hanretty, Chris, Benjamin E. Lauderdale, and Nick Vivyan. 2018. "Comparing Strategies for Estimating Constituency Opinion from National Survey Samples." *Political Science Research and Methods* 6 (3): 571–91.
- Kastellec, Jonathan P. 2018. "How Courts Structure State-Level Representation." *State Politics and Policy Quarterly* 18 (1): 27–60.
- Kastellec, Jonathan P., Jeffrey Lax, Michael Malecki, and Justin Phillips. 2015. "Polarizing the Electoral Connection: Partisan Representation in Supreme Court Confirmation Politics." *Journal of Politics* 77 (3): 787–804.
- Kastellec, Jonathan P., Jeffrey Lax, and Justin Phillips. 2010. "Public Opinion and Senate Confirmation of Supreme Court Nominees." *Journal of Politics* 72 (3): 767–84.
- Kiewiet de Jonge, Chad P., Gary Langer, and Sofi Sinozich. 2018. "Predicting State Presidential Election Results with National Tracking Polls and MRP." *Public Opinion Quarterly* 82 (3): 419–46.
- Krimmel, Katherine, and Kelly Rader. 2016. "Using Multilevel Regression and Poststratification on Historical Survey Data." Unpublished manuscript, August 27. <https://www.dropbox.com/s/r31fxicw7061amj/Historical%20MRP%20Paper%2016.07.20.pdf?dl=0>.
- Lax, Jeffrey, and Justin Phillips. 2009a. "Gay Rights in the States: Public Opinion and Policy Responsiveness." *American Political Science Review* 103 (3): 367–86.
- Lax, Jeffrey, and Justin Phillips. 2009b. "How Should We Estimate Public Opinion in the States?" *American Journal of Political Science* 53 (1): 107–21.
- Lax, Jeffrey, and Justin Phillips. 2012. "The Democratic Deficit in the States." *American Journal of Political Science* 56 (1): 148–66.
- Lax, Jeffrey, and Justin Phillips. 2013. "How Should We Estimate Subnational Opinion using MRP? Preliminary Findings and Recommendations." Working paper.
- Leemann, Lucas, and Fabio Wasserfallen. 2016. "The Democratic Effect of Direct Democracy." *American Political Science Review* 110 (4): 750–62.
- Leemann, Lucas, and Fabio Wasserfallen. 2017. "Extending the Use and Prediction Precision of Subnational Public Opinion Estimation." *American Journal of Political Science* 61 (4): 1003–22.
- Lewis, Daniel C., and Matthew L. Jacobsmeier. 2017. "Evaluating Policy Representation with Dynamic MRP Estimates: Direct Democracy and Same-Sex Relationship Policies in the United States." *State Politics and Policy Quarterly* 17 (4): 441–64.
- Lewis, Jeffrey B., and Drew A. Linzer. 2005. "Estimating Regression Models in Which the Dependent Variable Is Based on Estimates." *Political Analysis* 13 (4): 345–64.
- Linzer, Drew A. 2013. "Dynamic Bayesian Forecasting of Presidential Elections in the States." *Journal of the American Statistical Association* 108 (501): 124–34.
- Little, Roderick J., and Sonya Vartivarian. 2005. "Does Weighting for Non-response Increase the Variance of Survey Means?" *Survey Methodology* 31 (2): 161–68.
- Lowery, David, Virginia Gray, and Gregory Hager. 1989. "Public Opinion and Policy Change in the American States." *American Politics Research* 17 (1): 3–31.
- Lumley, Thomas S. 2010. *Complex Surveys: A Guide to Analysis Using R*. Hoboken, NJ: Wiley.
- Miller, Warren E., and Donald E. Stokes. 1963. "Constituency Influence in Congress." *American Political Science Review* 57 (1): 45–56.
- Movement Advancement Project. 2018. "Equality Maps: Dates of Laws and Policies." <http://www.lgbtmap.org/equality-maps>.
- Pacheco, Julianna. 2011. "Using National Surveys to Measure Dynamic US State Public Opinion: A Guideline for Scholars and an Application." *State Politics and Policy Quarterly* 11 (4): 415–39.
- Pacheco, Julianna. 2012. "The Social Contagion Model: Exploring the Role of Public Opinion on the Diffusion of Antismoking Legislation across the American States." *Journal of Politics* 74 (1): 187–202.
- Pacheco, Julianna. 2013. "The Thermostatic Model of Responsiveness in the American States." *State Politics and Policy Quarterly* 13 (3): 306–32.
- Pacheco, Julianna. 2014. "Measuring and Evaluating Changes in State Opinion across Eight Issues." *American Politics Research* 42 (6): 986–1009.
- Palus, Christine Kelleher. 2010. "Responsiveness in American Local Governments." *State and Local Government Review* 42 (2): 133–50.
- Park, David K., Andrew Gelman, and Joseph Bafumi. 2004. "Bayesian Multilevel Estimation with Poststratification: State-Level Estimates from National Polls." *Political Analysis* 12 (4): 375–85.
- Sances, Michael W., and Joshua D. Clinton. 2017. "New Policy, New Politics? The Effect of Medicaid Expansion on Public Support for the Affordable Care Act." Unpublished manuscript.
- Särndal, Carl-Erik, and Sixten Lundström. 2005. *Estimation in Surveys with Nonresponse*. Hoboken, NJ: Wiley.

- Selb, Peter, and Simon Munzert. 2011. "Estimating Constituency Preferences from Sparse Survey Data using Auxiliary Geographic Information." *Political Analysis* 19 (4): 455–70.
- 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 A* 178 (1): 1–28.
- Shor, Boris. 2017. "Ideology, Party and Opinion: Explaining Individual Legislator ACA Implementation Votes in the States." Unpublished manuscript, University of Houston.
- Tausanovitch, Chris, and Christopher Warshaw. 2013. "Measuring Constituent Policy Preferences in Congress, State Legislatures and Cities." *Journal of Politics* 75 (2): 330–42.
- Tausanovitch, Chris, and Christopher Warshaw. 2014. "Representation in Municipal Government." *American Political Science Review* 108 (3): 605–41.
- Treier, Shawn, and Simon Jackman. 2008. "Democracy as a Latent Variable." *American Journal of Political Science* 52 (1): 201–17.
- Warshaw, Christopher. 2016. "The Application of Big Data in Surveys to the Study of Elections, Public Opinion, and Representation." In R. Michael Alvarez, ed., *Computational Social Science: Discovery and Prediction*. New York: Cambridge University Press.
- Warshaw, Christopher, and Jonathan Rodden. 2012. "How Should We Measure District-Level Public Opinion on Individual Issues?" *Journal of Politics* 74 (1): 203–19.