

James Bisbee<sup>1</sup> and Jan Zilinsky<sup>2</sup>

<sup>1</sup>Center for Social Media and Politics, New York University, New York, NY, USA

<sup>2</sup>Technical University of Munich, Munich, Germany

## Abstract

The link between objective facts and politically relevant beliefs is an essential mechanism for democratic accountability. Yet the bulk of empirical work on this topic measures objective facts at whatever geographic units are readily available. We investigate the implications of these largely arbitrary choices for predicting individual-level opinions. We show that varying the geographic resolution—namely aggregating economic data to different geographic units—influences the strength of the relationship between economic evaluations and local economic conditions. Finding that unemployment claims are the best predictor of economic evaluations, especially when aggregated at the commuting zone or media market level, we underscore the importance of the modifiable areal unit problem. Our methods provide an example of how applied scholars might investigate the importance of geography in their own research going forward.

*Keywords:* modifiable areal unit problem, political geography, spatial politics, economic evaluations, partisanship, random forest, model selection, machine learning, ecological inference

## 1 Introduction

How are economic evaluations influenced by the unemployment rate? How is social trust influenced by ethnic diversity? How are concerns about crime related to the crime rate? Understanding the relationship between contextual phenomena and political opinions is central to social scientific research. Yet researchers often rely on the available geographic units at which contextual measures are aggregated with little attention paid to how this constraint influences their conclusions.

In this letter, we combine geographically rich data with machine learning methods to demonstrate that these choices carry nontrivial implications. Specifically, we show that the influence of “local” measures of the economy on economic evaluations varies substantially depending on the geographic unit at which we aggregate these contextual predictors. Substantively, in the face of a growing consensus in the literature arguing that politics is increasingly nationalized, our results emphasize the primacy of place in American politics.

In so doing, we highlight the continuing importance of the modifiable areal unit problem (MAUP) to political scientists.<sup>1</sup> The MAUP describes the statistical challenges associated with aggregating data from individual points of interest to geographic units, thereby compressing variation in the smaller units unless values were constant across them. The implications of using a particular measurement unit can be divided into two categories: the rigidity of borders and the salience of proximity. In contexts where geographic borders accurately demarcate differences in quantities of interest—that is, state-level policies, pork targeting a congressional district, and so on—the choice of which geographic unit to aggregate to is straightforward. Less understood is how the lived experiences of individuals are defined by these borders (but see Ansolabehere, Meredith,

*Political Analysis* (2023)  
vol. 31: 288–294  
DOI: 10.1017/pan.2021.50

**Published**  
22 February 2022

**Corresponding author**  
James Bisbee

**Edited by**  
Jeff Gill

© The Author(s) 2022. Published by Cambridge University Press on behalf of the Society for Political Methodology.

<sup>1</sup> The modifiable areal unit problem is related to, but distinct from the ecological inference fallacy (Robinson 1950), that is flawed inferences about individual-level relationships on the basis of aggregate data. See Section 4 in Supporting Information for a discussion.

and Snowberg 2014), yet these experiences are essential to accurately linking public opinion with local contexts.

We investigate the importance of the MAUP in the context of economic evaluations—an increasingly politicized dimension of opinion which presents a relatively hard test. We show that the decision to measure local economic factors at one geographic unit versus another matters for the empirical analysis of public opinion in terms of the overall model fit, the importance of contextual factors versus partisanship, and even the regression coefficients relating the two. We hope that, by underlining the degree to which these choices exert influence on the substantive conclusions drawn regarding a seminal dimension of American politics, our letter revitalizes the attention paid to this consequential decision in the data collection process, and stimulates innovations in the sources of data used by scholars to describe the microfoundations of politics.

## 2 Data and Methods

We investigate the degree to which an individual’s economic evaluation is predicted by contextual measures of the economy, where such measures are aggregated to different units. We do so by combining daily Gallup opinion data with local economic data based on tax returns from the Internal Revenue Service.<sup>2</sup>

The daily Gallup surveys randomly sample 1,000 American adults living across the United States, resulting in almost 1.7 million observations (respondents whose economic evaluations were elicited) for the period of our analysis (2008–2017). We examine our respondents’ assessment of the country’s economic conditions, where respondents can choose one of “poor,” “only fair,” “good,” or “excellent.” For each respondent we know their ZIP code of residence, allowing us to geolocate them with a high degree of accuracy.

We use an administrative data source—the federal U.S. tax authority—to obtain data at the ZIP code level on objective economic conditions. Our primary contextual measures of interest are the adjusted gross income (AGI) per return (logged thousands of dollars), unemployment compensation per return (logged thousands of dollars + 1), and the Gini coefficient. In addition, we control for the proportion of the population filing at each unit of aggregation. We provide a detailed description of these variables in the Supporting Information.<sup>3</sup>

Using crosswalk and shape files, we then calculate all our measures of interest for the most common geographic units available to researchers, summarized in Table 1. To match the place-based data with individual-level opinions, we use the ZIP code of each respondent to place them in the county where they live, their congressional district, their commuting zone, and so on.<sup>4</sup>

To evaluate the impact of the MAUP, we predict the evaluation of the economy  $y$  for a respondent  $i$  living in location  $j$  in year  $t$  using individual-level covariates  $\mathbf{X}_{it}$  (age, race, education, gender, marital status, self-reported income, and party ID), and contextual predictors  $\mathbf{G}_{jt}$  (AGI, income inequality, unemployment compensation, and proportion filing), along with year dummies  $\mathbb{1}_t$ .

$$y_{ijt} = f(\mathbf{X}_{it}, \mathbf{G}_{jt}, \mathbb{1}_t). \quad (1)$$

We are substantively interested in the impact of the unit at which we aggregate these contextual predictors  $\mathbf{G}_{jt}$  on three metrics: overall model fit, variable importance, and partial correlations.

To calculate the first two metrics of interest, we implement a random forest method, which relieves us of having to specify the correct functional form *a priori*. Overall model fit is calculated

<sup>2</sup> Replication materials are available at Bisbee and Zilinsky (2021).

<sup>3</sup> We underscore that because these are administrative data, they are less subject to the biases associated with the self-reported income asked on surveys.

<sup>4</sup> We recognize that the ZIP code is not the most precise unit but we are limited by data availability. A full description of the choices and challenges of aggregating the same measures to different geographic units is presented in the Supporting Information, Section 1.3.

**Table 1.** Size statistics for the data in 2016.

	Median population	Average population	Average area (sq mi)
Census region	39,022,232	35,627,715	393,582
State	4,425,976	6,287,244	69,456
Designated market area (DMA)	762,271	1,530,783	16,822
Congressional district	733,605	735,559	8,088
Commuting zone	121,438	452,895	5,003
County	25,660	102,137	1,127
Zip	4,334	11,069	88
Nation		323,127,513	3,540,144

Note: Measures do not include Alaska and Puerto Rico.

as the mean squared error (MSE) of the model's predictions, and variable importance is measured as the percent deterioration in MSE when information contained in a particular variable is removed via randomly reshuffling its values.<sup>5</sup> To estimate partial correlations, we model economic evaluations as a linear function of individual-level and geographic predictors via standard OLS.<sup>6</sup>

### 3 Does Geography Matter?

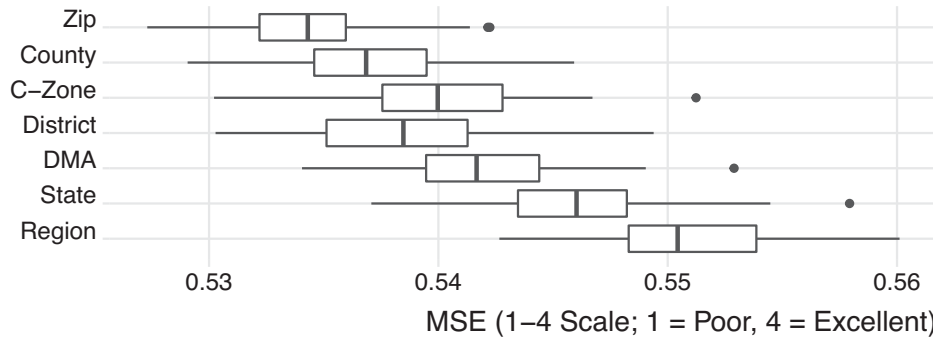
We begin by investigating how the choice of geography influences our ability to predict individuals' views of the economy. [Figure 1](#) plots the MSE of random forests that predict a respondent's view of the economy as a function of their individual-level covariates and contextual measures of the economy. These contextual measures are aggregated to different geographic units, ranging from the ZIP code to the Census subregion, indicated on the y-axis.

As the figure illustrates, the choice of the unit of aggregation matters for our ability to accurately predict the public's economic evaluations. However, while these differences are statistically significant, their substantive magnitude is small, corresponding to only 0.015 on a four point scale (mean = 1.86, standard deviation [SD] = 0.79 over the period of analysis), or a 2.7% increase in predictive accuracy when comparing the smallest and largest geographic units.

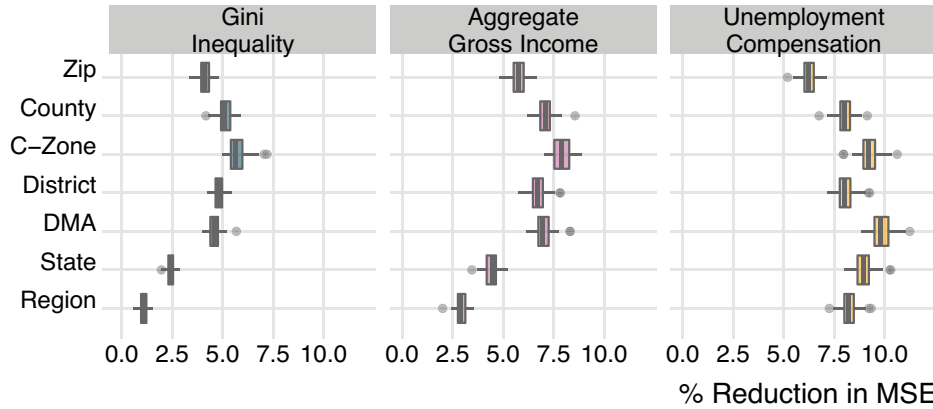
Just because these models perform better with contextual information aggregated to certain geographic units, does not necessarily mean that contextual predictors are more important in a substantive sense. To evaluate the impact of these choices on the predictive power of contextual data, we turn to permutation tests of variable importance. [Figure 2](#) plots the percent reduction in MSE associated with breaking the empirical relationship between AGI, income inequality, and unemployment compensation when aggregated to different geographic units.

Substantively, one might conclude that economic factors are unimportant when aggregated to the state or region, particularly for local inequality. But we find evidence that these variables matter most when aggregated to the commuting zone or—in the case of unemployment compensation—the designated market area, improving model accuracy by 5–10%. Furthermore, with the exception of the Congressional District, the relationship between the size of the geographic unit and the importance of the contextual variables aggregated within its borders is

- 5 We prefer the more computationally expensive permutation method given its insulation from bias when comparing continuous, categorical, and dichotomous predictors (Nicodemus et al. 2010), a benefit we illustrate with simulations in the Supporting Information, Section 3.
- 6 We include a variety of alternative specifications, with and without fixed effects, and a multilevel model, in our Supporting Information, Section 2.3.



**Figure 1.** Goodness of model fit: mean squared error (MSE, x-axis) by unit of geographic aggregation (y-axis).



**Figure 2.** Variable importance of income inequality (left), aggregate gross income (center), and unemployment compensation per return (right), aggregated to different geographic units (y-axes).

inverted U-shaped.<sup>7</sup> These patterns are consistent with the theory presented in Ansolabehere, Meredith, and Snowberg (2014) who argue that individuals choose information environments subject to a bias-variance trade-off.<sup>8</sup>

#### 4 Substantive Implications

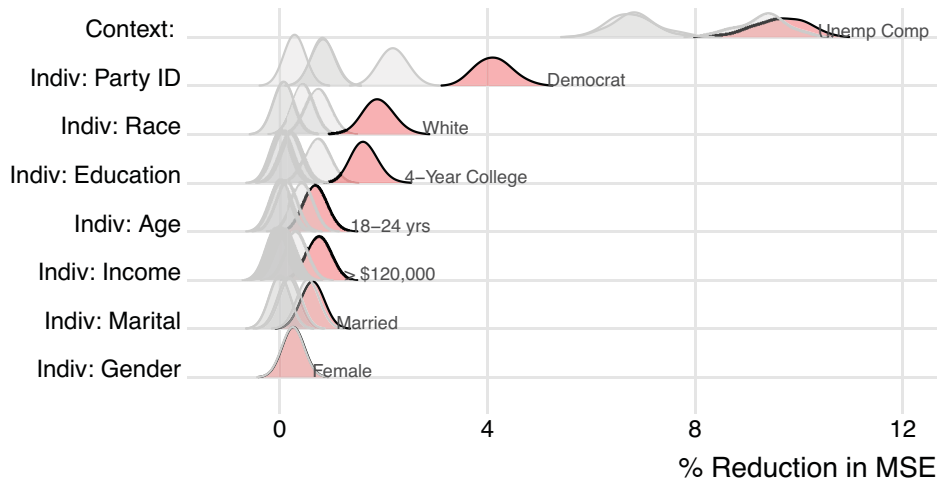
Thus far, we have shown that the choice of aggregation matters to both model fit, and for the importance of contextual-level predictors. However, are these differences large enough to change substantively important relationships in our data? We investigate this question in two ways.

First, we again rely on random forest permutation tests to compare the importance of our contextual measures to individual-level predictors. Figure 3 presents the variable importance of the top 5 most important predictors as densities where the contextual measures are aggregated at the commuting zone-level, highlighting that contextual measures are, in some cases, more than twice as important as the most prognostic individual-level covariates (4% reduction in MSE for Democrats versus an almost 9% reduction for unemployment claims). These patterns are attenuated when aggregating to larger units, the results for which are included in our Supporting Information.


Our second strategy for characterizing the substantive implications of these choices abandons random forests in favor of a simpler linear regression. Specifically, we estimate the partial correlation between AGI and positive views of the economy, controlling for individual-level characteristics

7 We suspect that the weaker importance for contextual measures aggregated to the Congressional District likely reflects gerrymandering, making these units particularly poor choices for measuring economic outcomes.

8 We re-analyze Ansolabehere, Meredith, and Snowberg (2014) in our Supporting Information, finding that their results about contextual unemployment rates are somewhat sensitive to the choice of the geographic unit of aggregation.



**Figure 3.** Variable importance of individual and contextual predictors, aggregated to the commuting zone, computed with bootstrapped permutation tests. The most important variable within each category is labeled and shaded.



**Figure 4.** Points are correlations between aggregate gross income (AGI) and positive views of the economy measured at the individual level, conditioning on respondents' demographic characteristics (including partisanship). Each row displays a coefficient from a separate model. The units at which we aggregate AGI are indicated on the y-axis.

and implementing year fixed effects.<sup>9</sup> We vary the geographic unit at which we aggregate AGI and present the coefficients along with two-standard-error bars in Figure 4.

We see again that the choice of the geographic unit influences the substantive conclusions one would draw about the relationship between an objective measure of local economic conditions and beliefs about the overall health of the American economy. We find that local income is associated with more positive evaluations of the economy when aggregated to smaller geographic units.<sup>10</sup> But as we measure AGI at larger geographic units, we observe estimates that seemingly suggest an insignificant association between economic conditions and evaluations of the economy. The choice of the unit of aggregation thus carries substantive implications when it comes to examining whether and how economic reality covaries with evaluations of the economy.<sup>11</sup>

<sup>9</sup> We divide AGI by the total number of tax returns filed to obtain a measure of average local income.

<sup>10</sup> A 1-SD increase in average local income is associated with a 0.027-point increase on the economic evaluations scale when aggregating to the ZIP code level.

<sup>11</sup> We present a variety of alternative specifications in our Supporting Information, Section 2.3.

## 5 Conclusion

The question of how individuals incorporate contextual information when forming political beliefs is of both theoretical and practical importance (Newman, Johnston, and Lown 2014). Substantively, democracy's normative appeal is predicated on the ability of individuals to perceive local welfare and adjust their political opinions accordingly. Methodologically, assessing the competing influence of objective facts and partisan motivated reasoning requires accurate measures of each.

Recent work on the impact of contextual variables on politics has suggested that proximate economic conditions have greater impact on public opinion than national economic outcomes (Bisgaard, Dinesen, and Sonderskov 2016; Newman 2020) but the existing work does not systematically investigate the sensitivity of effect sizes to different geographic units of aggregation. In this letter, we provide an evaluation of the choices researchers make when measuring objective facts. Focusing our investigation on economic evaluations—an increasingly politicized dimension of opinion (de Geus 2019)—we present a hard test of the importance of the MAUP in political science.

We combine machine learning tools with a rich dataset to show that the choice of geographic unit of aggregation has nontrivial consequences, demonstrating a monotonic decline in the predictive accuracy of a random forest as we aggregate to larger units, and significantly weaker correlation coefficients when estimating relationships using a linear regression. We also show that contextual measures of income, inequality, and unemployment are the most important predictors of an individual's assessment of the economy when aggregated to the individual's commuting zone. These results attenuate at smaller and larger units of aggregation, an empirical pattern consistent with the theory of “macro-economic voting” (Ansolabehere, Meredith, and Snowberg 2014) in which the optimal size of an individual's information environment is defined as a Goldilocks problem.

That we find meaningful differences in variable importance and model fit across the units of aggregation underscores the care required when predicting individual-level outcomes with contextual data. The tools we apply to this question can be used in other contexts to guide applied researchers when investigating the sensitivity of their results to these choices.

## Acknowledgments

We thank Neal Beck, Charlotte Cavaille, Pat Egan, Jared Finnegan, Gerda Hooijer, Sean Kates, Elif Kalaycioglu, Jonathan Nagler, Francesca Parente, Julia Payson, Abigail Vaughn, Mitch Watkins, and Ryan Weldzius for helpful comments.

## Data Availability Statement

Replication code for this article is available at (Bisbee and Zilinsky 2021)

## Supplementary Material

For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2021.50>.

## References

- Ansolabehere, S., M. Meredith, and E. Snowberg. 2014. “Macro-Economic Voting: Local Information and Micro-Perceptions of the Macro-Economy.” *Economics & Politics* 26 (3): 380–410.
- Bisbee, J., and J. Zilinsky. 2021. “Replication Data for: Geographic Boundaries and Local Economic Conditions Matter for Views of the Economy,” <https://doi.org/10.7910/DVN/C1XZKK>, Harvard Dataverse, V1.
- Bisgaard, M., P. T. Dinesen, and K. M. Sonderskov. 2016. “Reconsidering the Neighborhood Effect: Does Exposure to Residential Unemployment Influence Voters' Perceptions of the National Economy?” *The Journal of Politics* 78 (3): 719–732.
- de Geus, R. A. 2019. “When Partisan Identification and Economic Evaluations Conflict: A Closer Look at Conflicted Partisans in the United States.” *Social Science Quarterly* 100(5): 1638–1650.

- Newman, B. J. 2020. "Inequality Growth and Economic Policy Liberalism: An Updated Test of a Classic Theory." *The Journal of Politics* 82 (2): 765–770.
- Newman, B. J., C. D. Johnston, and P. L. Lown. 2014. "False Consciousness or Class Awareness? Local Income Inequality, Personal Economic Position, and Belief in American Meritocracy." *American Journal of Political Science* 59 (2): 326–340.
- Nicodemus, K. K., J. D. Malley, C. Strobl, and A. Ziegler. 2010. "The Behaviour of Random Forest Permutation-Based Variable Importance Measures Under Predictor Correlation." *BMC Bioinformatics* 11 (1): 1–13.
- Robinson, W. S. 1950. "Ecological Correlations and the Behavior of Individuals." *American Sociological Review* 15: 351–357.