UNIVERSITY OF CALIFORNIA

Los Angeles

Measuring and Modeling Partisan Geographic Segregation with Expected Kullback-Leibler Divergence

A thesis submitted in partial satisfaction

of the requirements for the degree

Master of Science in Statistics

by

Graham Paley Straus

2026

ABSTRACT OF THE THESIS

Measuring and Modeling Partisan Geographic Segregation

with Expected Kullback-Leibler Divergence

by

Graham Paley Straus

Master of Science in Statistics

University of California, Los Angeles, 2026

Professor Mark S. Handcock, Chair

Democrats often live with minimal exposure to Republicans and Republicans often live with minimal exposure to Democrats. Leading accounts of partisan segregation document intense and rising segregation both across and within counties. Existing approaches for measuring partisan segregation either drop nonpartisans altogether or model their partisanship. These approaches also use single scale measures of segregation, like the index of dissimilarity. In this thesis, I measure partisan segregation with national voter registration records and a multiscale approach—Expected Kullback-Leibler Divergence—finding stable and relatively low levels of within-county segregation. I further describe patterns in segregation with a spatiotemporal conditional autoregressive model.

The thesis of Graham Paley Straus is approved.

Jeffrey B. Lewis

Frederic Roland Paik Schoenberg

Mark S. Handcock, Committee Chair

University of California, Los Angeles

2026

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction (Literature Review)

Partisans are segregated (Brown and Enos 2021). Democrats are isolated from Republicans and Republicans are isolated from Democrats. In large cities Democrats live with minimal exposure to Republicans and the opposite is true for vast swaths of rural America. Both academics (Kaplan, Spenkuch, and Sullivan 2022; Bishop 2009; Sussell 2013) and journalists (Badger, Quealy, and Katz 2021; Kaysen and Singer 2024) have taken a strong interest in this trend because segregation between groups—even partisan groups—is associated with negative outcomes like prejudice and conflict (Massey and Denton 1988; Enos 2017).

Moving is, surprisingly, not one of the main drivers of partisan segregation. Two recent studies confirm this, one that pairs survey responses on moving intentions with data on where respondents actually choose to live (Mummolo and Nall 2017) and one that decomposes partisan segregation into components: generational turnover, party switching, and residential mobility. Brown et al. (2025) write that rising geographic polarization "is less about Democrats and Republicans fleeing from each other and more about macro-level forces shaping the types of people each party represents and the demographic composition of the U.S. electorate." Survey experimental evidence shows that partisans review a home more favorably when they know the neighborhood shares their own partisan identity (Gimpel and Hui 2015) but the concerns that are top of mind when people move (e.g., schools and jobs) are not about partisanship directly but rather reflect factors correlated with partisanship (Ihlanfeldt and Yang 2024).

Why is partisan segregation an issue? One reason scholars offer is that partisan segregation leads to electoral bias (Chen and Rodden 2013), skewing representation in state and federal legislatures and exacerbating discrepancies between the popular vote and Electoral College (Stromberg 2008). Another consequence from intergroup contact theory is that contact should make partisans see the outgroup more favorably. This may be true for other identity groups (e.g., Massey and Denton 1988), but it is unclear if the same is true for partisanship. Studies that test contact effects have found that the subject of discussion matters for efficacy (Santoro and Broockman 2022) and that intergroup contact cannot change strong partisans' minds (Thomsen and Thomsen 2023). However, living close to out-partisans need not imply greater contact with the out-group. Abrams and Fiorina (2012) present the critique of the "Big Sort" hypothesis that people do not spend that much time interacting with their neighbors. Neighborhoods do not fully reflect the social bonds that matter most like friends, family, and coworkers. Furthermore, to the extent scholars have investigated the link between residential partisan segregation and affective polarization, there is no association (Tilley and Hobolt 2025).

A first order question concerns how to measure the segregation of partisans. Wrapped up in this question is how to classify partisans and how we measure spatial segregation. In this thesis I use a multiscale measure of segregation that better accounts for how people experience the area in which they live. I also include nonpartisans as their own category in modeling partisan segregation since in recent years nearly a third of people are registered without party affiliation. I draw on granular, individual-level voter registration data to show that partisan segregation did not increase between 2014 and 2020. Finally, I use a spatiotemporal conditional autoregressive model to delve into which areas show increasing or decreasing segregation.

2 Within-County Partisan Segregation

Segregation happens across and within geographies. In other words, different counties show different baseline rates of partisanship, but even within a given county partisans reside apart from one another. Consider, for example, Orange County (Figure 1). The cities, like Santa Ana and Anaheim, are more Democratic, and the areas outside the city are more Republican. When describing segregation across counties researchers are not speaking to the fact that blue clusters near blue and red clusters near red even *within* a given county. Orange County trended democratic between 2014 and 2020—it went from 43.8% Democratic to 52.3% Democratic according to registration data. And crucially, these changes happened across the whole county, not just in the already Democratic pockets, making for lower within-county measures of partisan segregation.

Traditionally scholars have used the index of dissimilarity (Massey and Denton 1988; Brown et al. 2025) to calculate a measure of unevenness, or the average difference between each micro-region and the macro-region; for instance, precincts within districts or tracts within counties. Formally, for $\{D, R\}$ = {n Democrats, n Republicans} countywide, $\{D_t, R_t\}$ = {n Democrats in tract t, n Republicans in tract t}, and T tracts; the index of dissimilarity (ID) is:

$$ID = \frac{1}{2}\sum_{t=1}^{T}\left|\frac{D_t}{D} - \frac{R_t}{R}\right| \tag{1}$$

There are also multigroup extensions of the index of dissimilarity like Theil's H (Reardon and Firebaugh 2002) but most of the existing literature on partisanship just uses Democrats and Republicans. One of the benefits of the index of dissimilarity is its ease of interpretation. For example, in Orange County the index of dissimilarity ranges between .20 and .22 between 2014 and 2020 meaning around 20% of one group would need to move to achieve evenness throughout the county.

Figure 1: Orange County Democratic Share of Registrants (left) and Within-County Partisan Segregation (right)

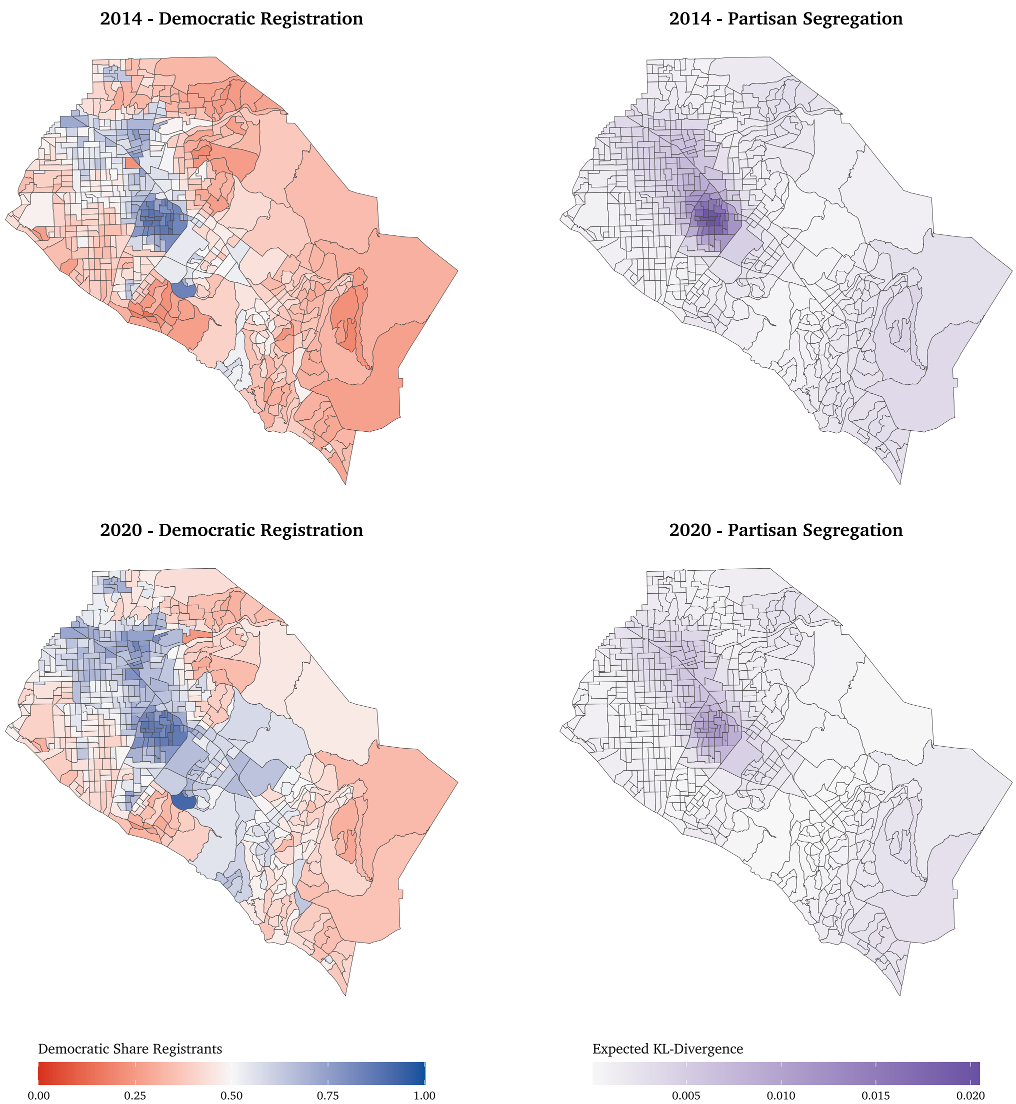

Note: L2 Voter registration data from Orange County, California. Shading at the tract level.

2.1 Segregation Through the Multiscalar Lens

As Olteanu, Randon-Furling, and Clark (2019) point out, these single scale indices miss many of the most interesting dynamics of how people experience place. No person only experiences their precinct or tract. Rather, people experience distinct parts of the city in some smooth fashion. Observing this, Olteanu, Randon-Furling, and Clark (2019) introduce a multiscale measure, the distortion coefficient, and to build intuition they write: "Consider a city, the population of which comprises k groups. Imagine that an individual explores the city from her home, looking first at her direct neighbors and then at the building next door, and so on, in an ever-expanding visual scale. She will gradually meet the whole city, but at every step, the proportions of the groups in the population that she has encountered so far produce a sequence that yields detailed information on the starting point, in multiscalar relationship with the whole city."

Further developing the multiscale measure of within-county segregation, Handcock et al. (2019) use expected Kullback-Leibler divergence (EKL) to quantify within-county segregation. The formal definition is as follows: Let j be the index of the unit a randomly chosen person from the population is from ($j \in U$) and let $d_{KL}(q^{i,1:\tilde{u}_{ij}}|r)$ be the Kullback-Leibler divergence between the distribution of the population in the units at least as close as unit j to unit i. Here $\tilde{u}_{ij}$ is the number of units closer to i than j. EKL is the expected value of $d_{KL}(q^{i,1:\tilde{u}_{ij}}|r)$ (over the random choice of j). There is an alternative index specification using random walks, considering an individual aware of regions in proportion to the probability she arrives there on a Brownian motion on the plane. The intuition is the same as in Olteanu, Randon-Furling, and Clark (2019), and is an improvement over single scale indices of dissimilarity or unevenness. The nature of this measure makes for smooth values of EKL since a person will encounter their immediate neighbors most often and the individuals they live farthest from least often.

Figure 2: Orange County EKL

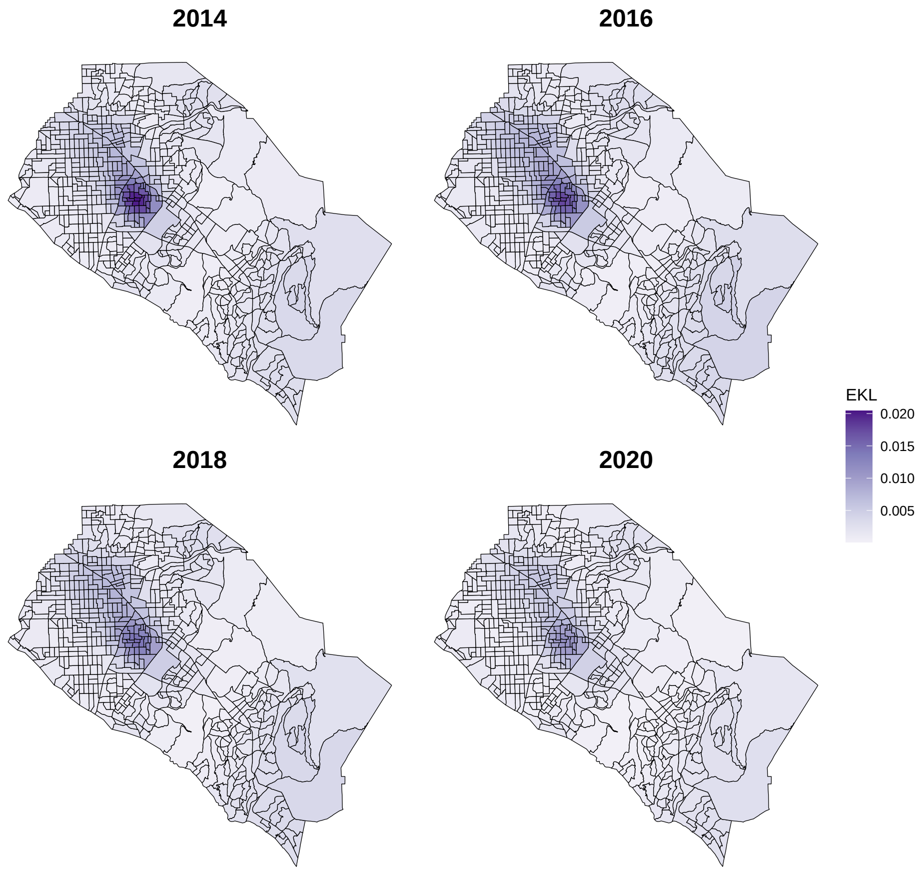

Note: L2 Voter registration data from Orange County, California; EKL shows within-county segregation of registered Democrats, Republicans, and nonpartisans.

Nearly 30% of people in voter files, even in states with registration partisanship, are listed without a party affiliation. This share has increased in the last decade. This presents a problem for using the index of dissimilarity to measure within-county partisan segregation since all independents or nonpartisans must have modeled partisanship or be dropped from the analysis. See Appendix Figure A.5 for a depiction of how self-identified partisanship aligns with voter file registration-based partisanship. I use three groups based on voter file partisanship: Democrats, Republicans, and nonpartisans. An overwhelming majority of independents do lean towards one party or the other (Atske 2019), however, over 70% of self-identified independents appear in voter files with no partisan affiliation. Overall it is more faithful to the granular registration data from which measures of segregation are derived to include the nonpartisans as their own group. I prefer this to dropping them or using modeled partisanship which itself is built off neighborhood characteristics.

2.2 Longitudinal Estimates of Segregation

There are lower average values of EKL (lower within-county partisan segregation) in more recent years. Figure 2 demonstrates this descriptive pattern in Orange County, CA. The dark purple area around the Anaheim region lightens between 2014 and 2020—tracts in this region become more like the county as a whole over this period. This is not just a comparison of individual tracts to the overall county. Rather, EKL considers how individuals in a given tract will experience their whole environment by assuming they spend the most time in the closest neighboring areas and progressively less time in areas farther away from home.

To test whether the pattern I observe in Orange County is true in the broader US, I use national voter registration files. These files come from the data vendor L2. L2 aggregates state registration records. In 30 states counties must report party affiliation for registered voters and these are the states I analyze below. Going county by county I calculate EKL for each tract in a given county-year. Below, in Figure 3, I show kernel density estimates of the tract-level EKL scores broken out by year. The average stays fairly constant between 2014 to 2020, decreasing slightly between 2014 and 2020. Table 1 shows, for the 1,224 counties in my sample, the annual averages of these scores. Table 1 also shows averages just among counties that in the aggregate became more partisan in this period.

In some respects, the findings on within-county segregation align with existing work, specifically from Brown et al. (2025). The same counties that have high segregation indices according to their approach (using the index of dissimilarity and two commercial voter files) have higher EKL scores. The national map in Figure 4 shades according to county-level averages of segregation and correlates highly with the corresponding figure in Brown et al. (2025). However, my year-by-year findings differ from existing work. The over-time estimates in Figure 3 and Table 1 differ from recent work that stresses the extreme and increasing isolation of partisans from each other. Accounting for nonpartisans—who make up 30% of registrants in national voter records—as their own category and using the multiscale EKL to score segregation reveals that within-county segregation is not increasing and may even be decreasing. To test this more rigorously I proceed by modeling segregation while accounting for spatial dependency.

Figure 3: Distribution of Tract-level Within-County Segregation (EKL) 2014-2020

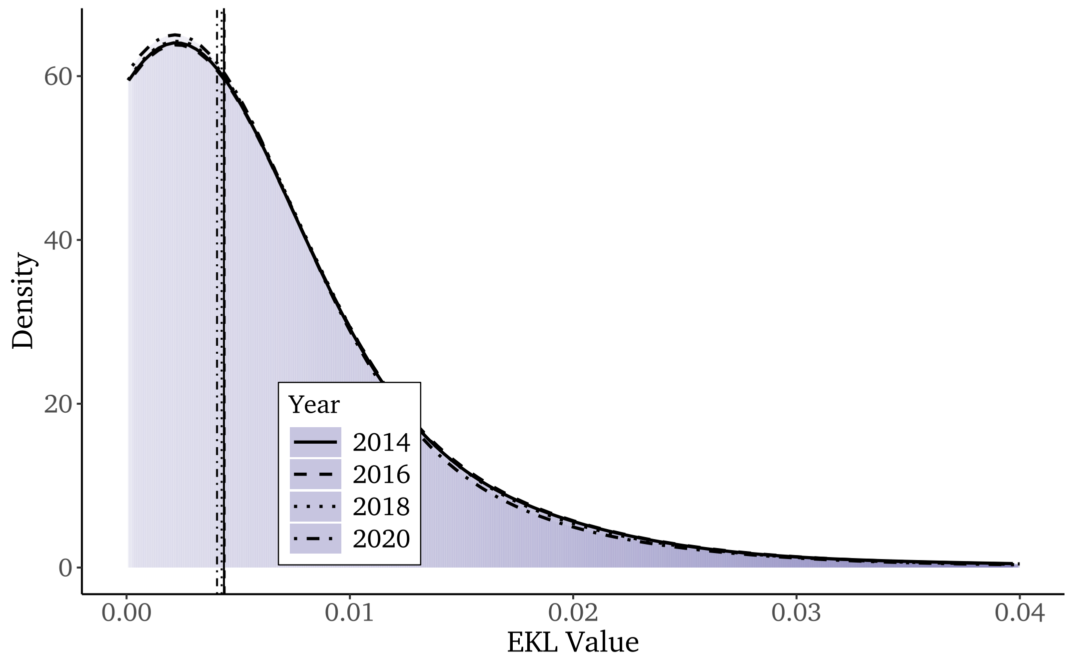

Note: L2 Voter registration data from 1,224 counties in states with registration partisanship; kernel density estimates are weighted by counts of registered voters in a given tract-year and use a Gaussian kernel with bandwidth of 0.005.

Table 1: EKL Within-County Segregation Averages by Year

Year	Weighted Mean	Dem Trending Counties	Rep Trending Counties
2014	0.00436	0.00462	0.00389
2016	0.00439	0.00476	0.00375
2018	0.00426	0.00453	0.00377
2020	0.00405	0.00425	0.00370

Note: L2 Voter registration data from 1,224 counties in states with registration partisanship. Dem and Rep trending counties indicate the 300 and 924 counties that respectively became more Democratic or Republican between 2014 and 2020.

Figure 4: National Map and Population-Based Cartogram Showing County Averages of Expected Kullback-Leibler Divergence

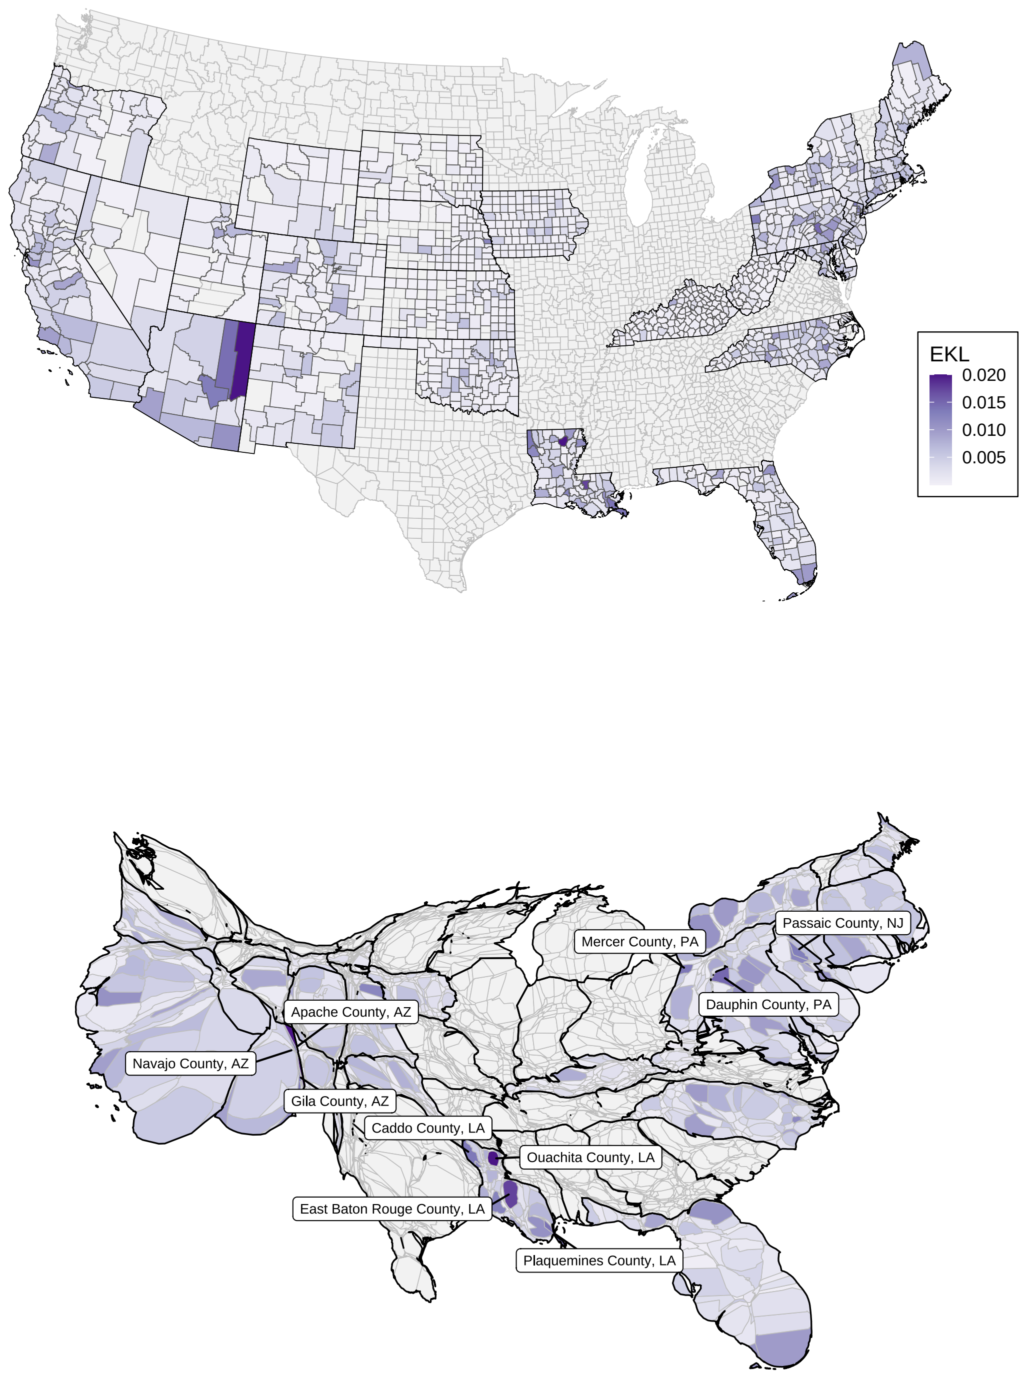

Note: Each county is shaded by its (population-weighted) average EKL. Data from 2018 L2 voter registration snapshot. The 10 counties with the highest average partisan segregation across tracts are labeled in the cartogram which expands county boundaries according to population.

3 Conditional Autoregressive Model of EKL

Partisan segregation is inherently a spatial phenomenon, and measurements taken across neighboring geographic units are not independent. By design, measurements for neighboring tracts using Expected Kullback-Leibler Divergence are not independent. Census tracts that share borders often exhibit similar demographic, economic, and political characteristics due to shared local contexts, housing markets, and social networks. Ignoring this spatial dependence when modeling segregation measures can lead to inefficient estimates and misleading inferences. Moreover, when observations are collected over multiple time periods, temporal dependence arises because the same populations and communities are observed repeatedly, and political sorting processes unfold gradually over time. A principled modeling approach must therefore account for both spatial and temporal autocorrelation simultaneously.

The model below provides a natural framework for decomposing spatio-temporal variation in the segregation measure into interpretable components. By separating the overall spatial pattern (common across all time periods) from the overall temporal trend (common across all units), I can assess whether partisan segregation exhibits stable geographic patterns, whether there are region-wide temporal shifts, or both. This ANOVA-style decomposition is useful for the application because I am interested in understanding not just where segregation is high or low, but also how segregation levels have evolved over the 2014–2020 period across counties.

The model employs conditional autoregressive (CAR) priors for both the spatial and temporal random effects, which provide a flexible yet parsimonious approach to modeling dependence structures in areal data. Specifically, I use the Leroux, Lei, and Breslow (2000) CAR prior, which includes a spatial (and temporal) dependence parameter ρ that controls the degree of smoothing. When $\rho = 1$, the model implies strong autocorrelation where neighboring units are tightly linked; when $\rho = 0$, the random effects are independent. Such models with $\rho = 1$ are sometimes referred to as "intrinsic" CAR models which assume stationarity of the differences between neighboring random effects rather than stationarity of the random field itself (Besag, York, and Mollié 1991). By estimating these parameters from the data, the model adaptively determines the appropriate level of spatial and temporal smoothing rather than imposing a fixed dependence structure. This flexibility is valuable given uncertainty about the precise nature of spatial and temporal dependencies in partisan segregation patterns.

3.1 Model Specification

For areal units (tracts) $u = 1, \ldots, U$ and time periods $t = 1, \ldots, N$:

$$Y_{ut} \sim \text{Normal}(\mu_{ut}, \nu^2)$$

$$\mu_{ut} = \mathbf{x}_{ut}^{\top}\boldsymbol{\beta} + \phi_u + \delta_t$$

where $\mathbf{x}_{ut}$ is a vector of covariates for unit u at time t with corresponding regression coefficients $\boldsymbol{\beta}$, ϕ_u is a spatial random effect for unit u, and δ_t is a temporal random effect for time period t. Since EKL is strictly positive and right-skewed, I model its natural logarithm, which is approximately normally distributed (See Appendix Figure A.1). Y_{ut} is $\log(\text{EKL})$ for tract u in time t.

3.1.1 Prior Specifications

The spatial random effects follow a Leroux CAR prior:

$$\phi_u \mid \phi_{-u}, W \sim \text{Normal}\left(\frac{\rho_S \sum_{j=1}^{U} w_{uj}\phi_j}{\rho_S \sum_{j=1}^{U} w_{uj} + 1 - \rho_S}, \frac{\tau_S^2}{\rho_S \sum_{j=1}^{U} w_{uj} + 1 - \rho_S}\right)$$

where W is a $U \times U$ binary spatial adjacency matrix with $w_{uj} = 1$ if units u and j share a common border and $w_{uj} = 0$ otherwise.

The temporal random effects follow a Leroux CAR prior:

$$\delta_t \mid \delta_{-t}, D \sim \text{Normal}\left(\frac{\rho_T \sum_{s=1}^{N} d_{ts}\delta_s}{\rho_T \sum_{s=1}^{N} d_{ts} + 1 - \rho_T}, \frac{\tau_T^2}{\rho_T \sum_{s=1}^{N} d_{ts} + 1 - \rho_T}\right)$$

where D is a $N \times N$ binary temporal adjacency matrix with $d_{ts} = 1$ if $|s - t| = 1$ and $d_{ts} = 0$ otherwise.

3.1.2 Hyperpriors

$$\begin{aligned}
\boldsymbol{\beta} &\sim \text{Normal}(\boldsymbol{\mu}_\beta, \boldsymbol{\Sigma}_\beta) \\
\tau_S^2, \tau_T^2 &\sim \text{Inverse-Gamma}(a, b) \\
\rho_S, \rho_T &\sim \text{Uniform}(0, 1) \\
\nu^2 &\sim \text{Inverse-Gamma}(a, b)
\end{aligned}$$

All random effects are mean-centered: $\sum_{u=1}^{U} \phi_u = 0$ and $\sum_{t=1}^{N} \delta_t = 0$.

I fit these models with `CARBayesST`, software that evaluates spatio-temporal Bayesian models with Markov chain Monte Carlo simulation (Lee, Rushworth, and Napier 2018). In the application below, it turns out that models with only an intercept term, such that $\mathbf{x}_{ut}^\top \boldsymbol{\beta} = \beta_0$, do just as well as models with covariates.

3.2 Fitting the Model to Orange County Data

As a first step I fit the model described above to longitudinal data from Orange County. I compare a version of the model with only the random effects to a version with tract-level Census covariates (Schroeder et al. 2025). These covariates are racial demographics, the fraction of individuals over 60 years old, the fraction of foreign born residents, log of the median household income, and the fraction of residents who have college-level education or above.

I compare the intercept-only model to the model with all covariates using three Bayesian model selection criteria. The Deviance Information Criterion (DIC) penalizes the posterior mean deviance by an estimated effective number of parameters. The Watanabe-Akaike Information Criterion (WAIC) improves on DIC by averaging over the full posterior distribution rather than conditioning on a point estimate, with a variance-based penalty. The Log Marginal Predictive Likelihood (LMPL) evaluates leave-one-out predictive performance via the conditional predictive ordinates. Lower DIC and WAIC and higher LMPL indicate better fit. The intercept-only model has lower WAIC and DIC as well as a greater log marginal predictive likelihood (Table 2). This suggests that partisan segregation patterns are not reducible to simple demographic components beyond spatial structure. Intra-county partisan sorting is a spatial phenomena and isn't just a byproduct of racial segregation or income stratification for example, even though those demographics are correlated with partisanship.

Figure 5: Observed and Fitted Values of Log EKL in Orange County in 2014 and 2020

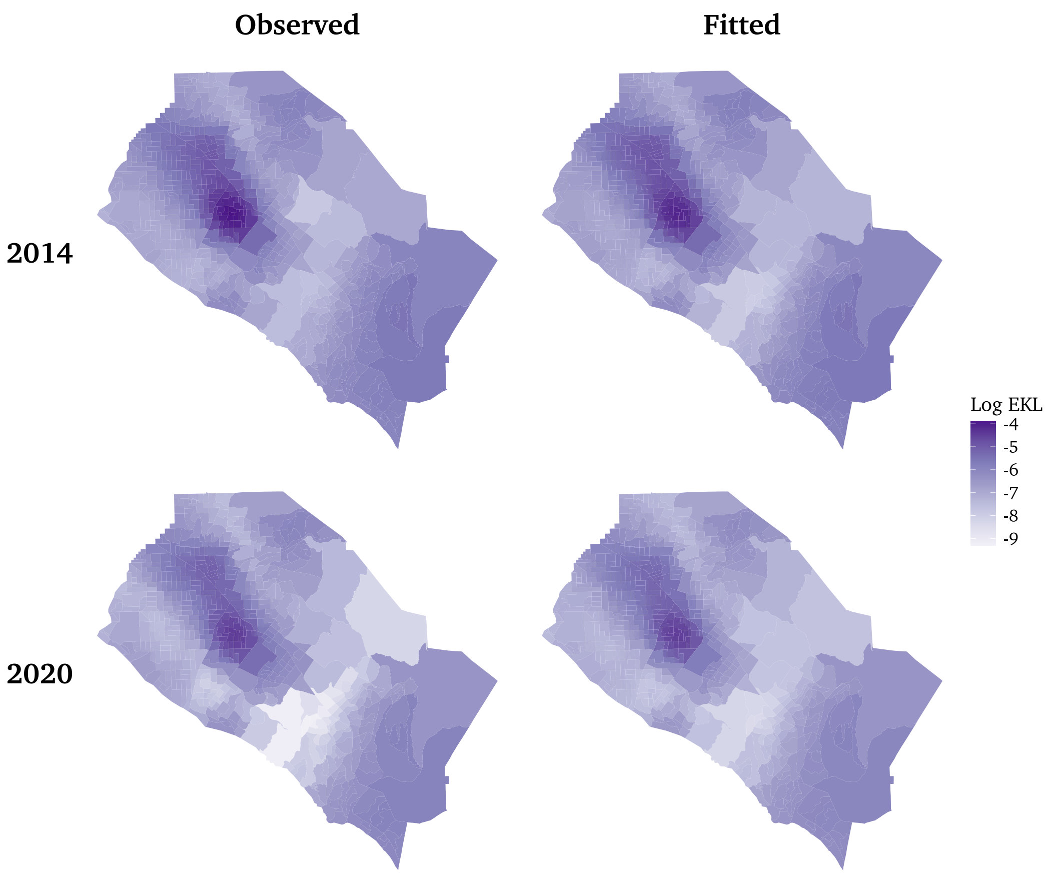

Table 2: Model Comparison (Orange County) — Intercept-Only vs. Covariates

Model	DIC	WAIC	LMPL
Intercept only	-814.41	-788.34	375.65
With covariates	-806.91	-779.06	370.38

Table 3: Parameter Estimates from Fitting the Intercept-Only Model

Parameter	Mean	2.5%	97.5%
β_0	-6.199	-6.207	-6.192
ν^2	0.033	0.031	0.035
τ_S^2	0.252	0.222	0.286
τ_T^2	0.044	0.012	0.142
ρ_S	0.997	0.990	1.000
ρ_T	0.428	0.022	0.923

Inference was conducted via MCMC with 3 parallel chains, each run for 50,000 iterations with a burn-in of 10,000 and thinning factor of 10, yielding 12,000 posterior samples (4,000 per chain). Trace plots are in Appendix Figure A.4.

3.3 Temporal Random Effects

To examine whether temporal patterns in segregation vary across counties, I fit separate intercept-only models with space-time interaction to each of California's 15 most populous counties (those with more than 100 census tracts). For each county, I extract the posterior distribution of the temporal random effects δ_t for $t \in \{2014, 2016, 2018, 2020\}$. Because each model is fit independently with its own intercept, the δ_t estimates represent within-county deviations from that county's baseline level of segregation rather than comparable absolute levels across counties.

Figure 6 displays the posterior means and 95% credible intervals for δ_t across all 15 counties, with four counties highlighted: Los Angeles, Orange, San Bernardino, and San Mateo. The results reveal substantial heterogeneity in temporal trajectories. Most counties show a stable or slightly negative pattern over time, consistent with the descriptive finding that within-county segregation declined slightly over this period. However, the magnitude and timing of changes vary considerably. San Mateo County exhibits a sharp decline in δ_t by 2020, while San Bernardino increases from 2014 to 2016 and then remains relatively stable throughout. Los Angeles and Orange Counties show moderate downward trajectories. The wide credible intervals for ρ_T in Table 3 suggest that the strength of temporal autocorrelation is not precisely identified with only four time points, though the spatial dependence parameter is estimated near 1 across all models, indicating the known strong spatial dependence.

Figure 6: δ_t From California's 15 Most Populous Counties

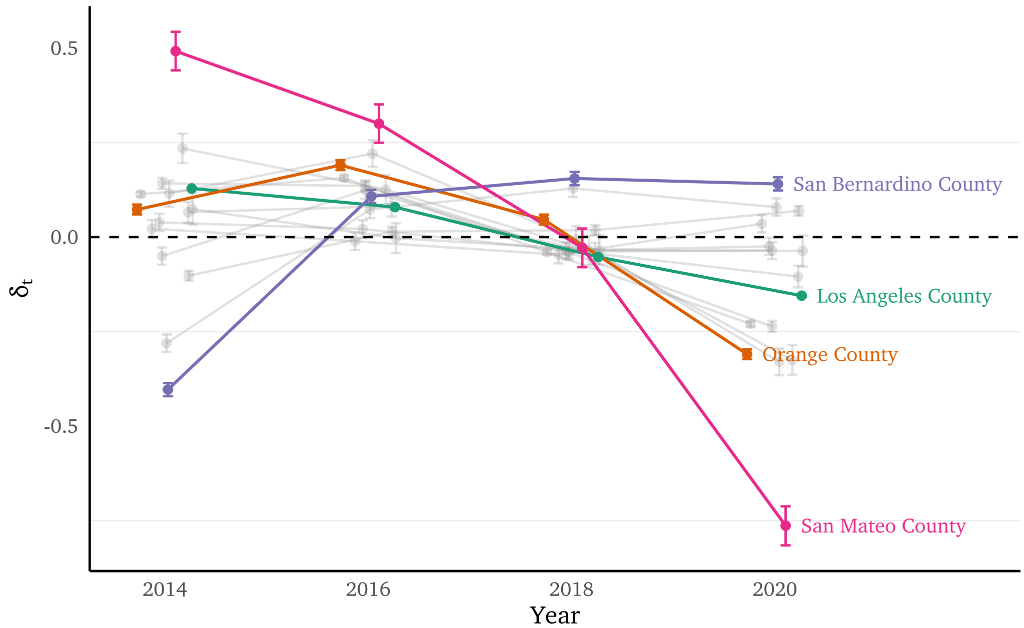

4 Discussion

This thesis makes two contributions: one to the study of partisan geography and one to the measurement and modeling of segregation. On the substantive side, the central finding is that within-county partisan segregation is stable or slightly declining between 2014 and 2020. This contrasts with the prevailing narrative that partisan sorting is extreme and increasing. The difference arises from two methodological choices. First, I include nonpartisans—who constitute nearly 30% of registrants in states with party registration—as their own category rather than dropping them or imputing their partisanship. Second, I use Expected Kullback-Leibler Divergence, a multiscale measure that captures how individuals actually experience their geographic environment rather than comparing fixed administrative units at a single scale. With these choices the picture that emerges is one of modest and stable within-county segregation.

To put the magnitude of partisan segregation in context, Figure 7 compares within-county partisan EKL to within-county racial EKL (computed with groups as White, Hispanic, and other). Racial segregation averages 5.9 times the magnitude of partisan segregation. While the two are positively correlated, the intercept-only CAR model's superiority over the covariate model suggests that partisan sorting is not simply a byproduct of racial or socioeconomic residential patterns. Partisan segregation has its own spatial structure that persists after accounting for where racial groups, income levels, and education levels concentrate.

Figure 7: Racial Segregation Correlates with Partisan Segregation but is Much More Severe

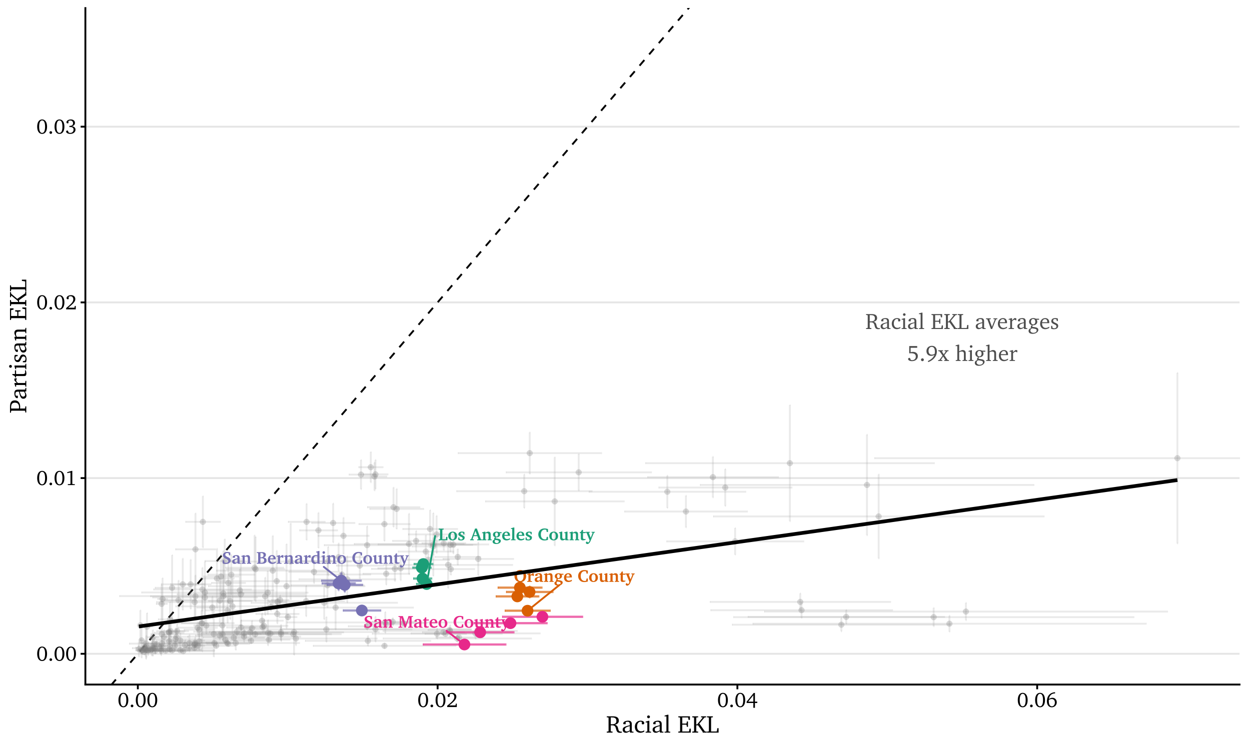

Note: Each point represents a county in a given year. The x-axis value corresponds to the county's average racial EKL (with groups as White, Hispanic, and other) and the y-axis value corresponds to the county's average partisan EKL (with groups as Democrats, Republicans, and nonpartisans). The dashed line is the 45 degree line which corresponds to equal magnitude estimates for racial and partisan segregation.

If partisan segregation were consequential for democratic health, it should be associated with affective polarization—the tendency for partisans to dislike the opposing party. Figure 8 plots county-level EKL against county-level affective polarization estimates from Warner (2025) and finds no clear association. Larger, more populous counties have more partisan segregation, but more segregated counties do not have more affective polarization. This is consistent with Tilley and Hobolt (2025), who find no link between residential partisan homogeneity and affective polarization in Britain. It may be that neighborhoods, as Abrams and Fiorina (2012) argue, are simply not the locus of the social interactions that drive partisan animosity. Friends, family, coworkers, and media consumption may matter far more than who lives on one's block.

Figure 8: Correlation Between Within-County Segregation and County-Level Affective Polarization

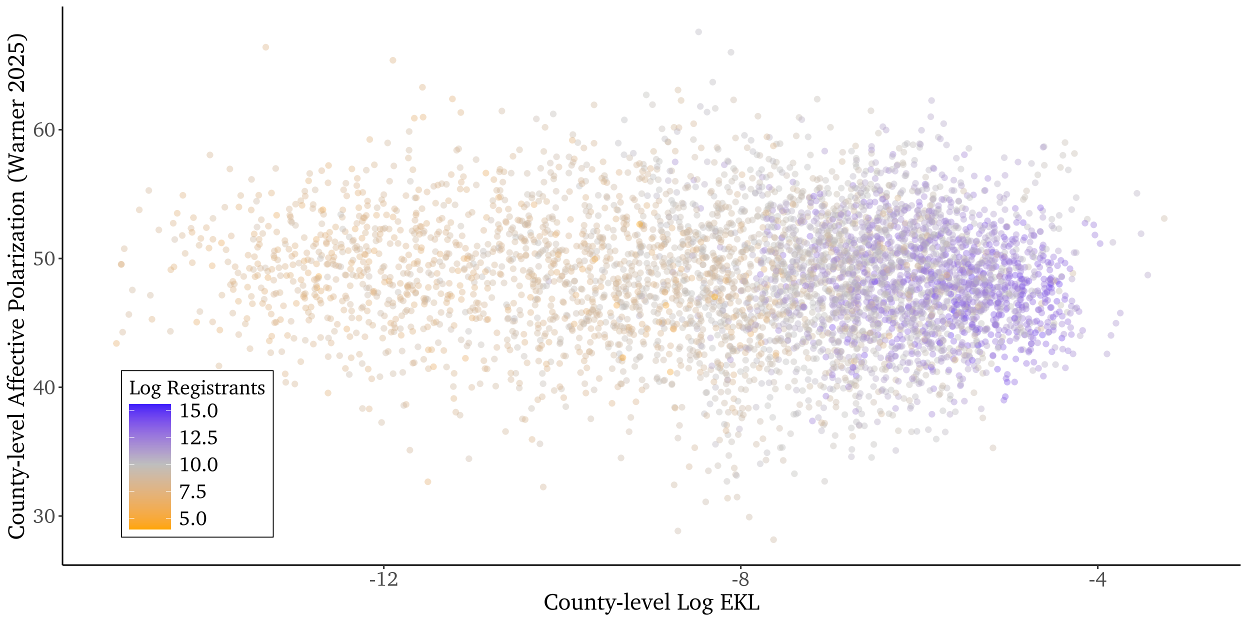

Note: L2 Voter registration data for EKL; Warner (2025) for affective polarization estimates. Higher value of log county-level EKL means more unevenness in where partisans live within the county. The shading in color corresponds to population with orange representing less populous and blue representing more populous counties. Figure uses data from all years 2014-2020.

On the methodological side, the spatiotemporal CAR model proves well-suited to this application. The spatial dependence parameter is estimated near 1 across all county models, confirming that partisan segregation exhibits strong spatial autocorrelation. With only four snapshots in time, the temporal dependence parameter is less precisely estimated. The heterogeneity in temporal trajectories across California's 15 largest counties (Figure 6) underscores the value of county-specific models; aggregate trends mask meaningful local variation.

Looking forward, there are several avenues for extending this work. The most immediate is incorporating additional voter registration snapshots, which would expand the number of time points and substantially improve the estimation of the temporal dependence parameter. The work in this thesis has focused on *intra*-county segregation, but additional work can do more to test national trends by considering how subunits differ from the partisan composition of the country as a whole.

Appendix

TABLE OF CONTENTS

A.1 Goodness of Fit Tests for Expected Kullback-Leibler Divergence

Figure A.1: Goodness of Fit with Data from Orange County, CA.

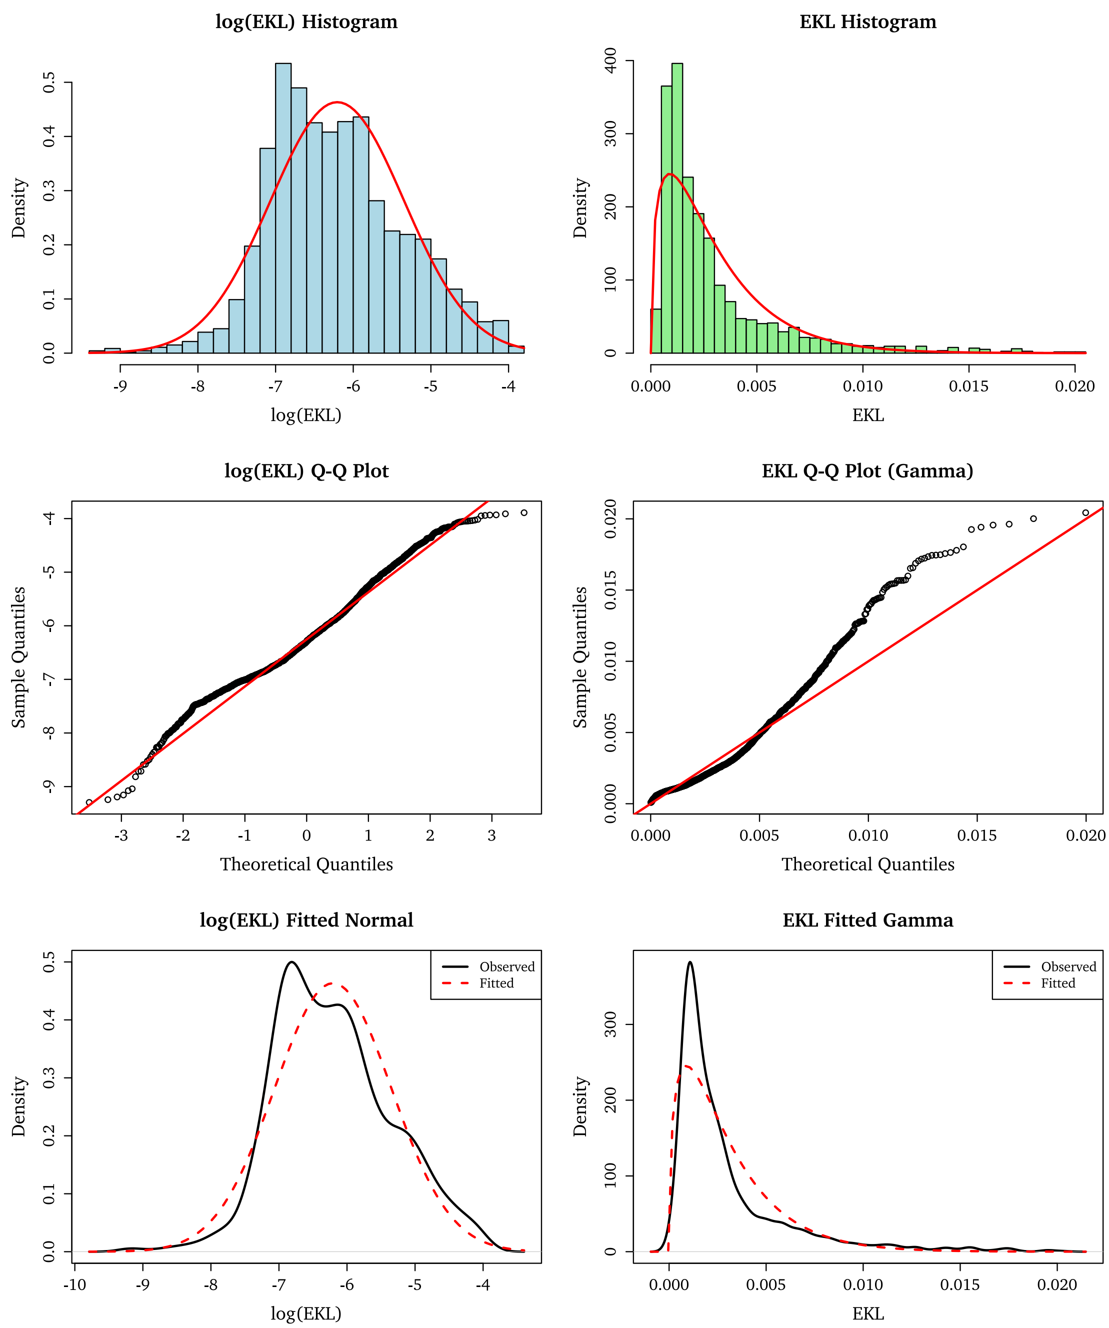

A.2 Across-County and Across-Precinct Segregation Decreased 2020-2024

In (Brown et al. 2025), the authors use Catalist and TargetSmart files to calculate the two-party Democratic share of registrants at various levels of geography ranging from counties down to neighborhoods. They show that the population-weighted standard deviation of Democratic share increases from 2008 to 2020, putting more mass in the tails (areas with more homogeneous partisan makeup). Surprisingly, this did not continue in 2024. The figure below shows population-weighted density plots for county-level and precinct-level election returns in 2020 and 2024.

Figure A.2: Density plots for 2020 and 2024 election results at the county and precinct level

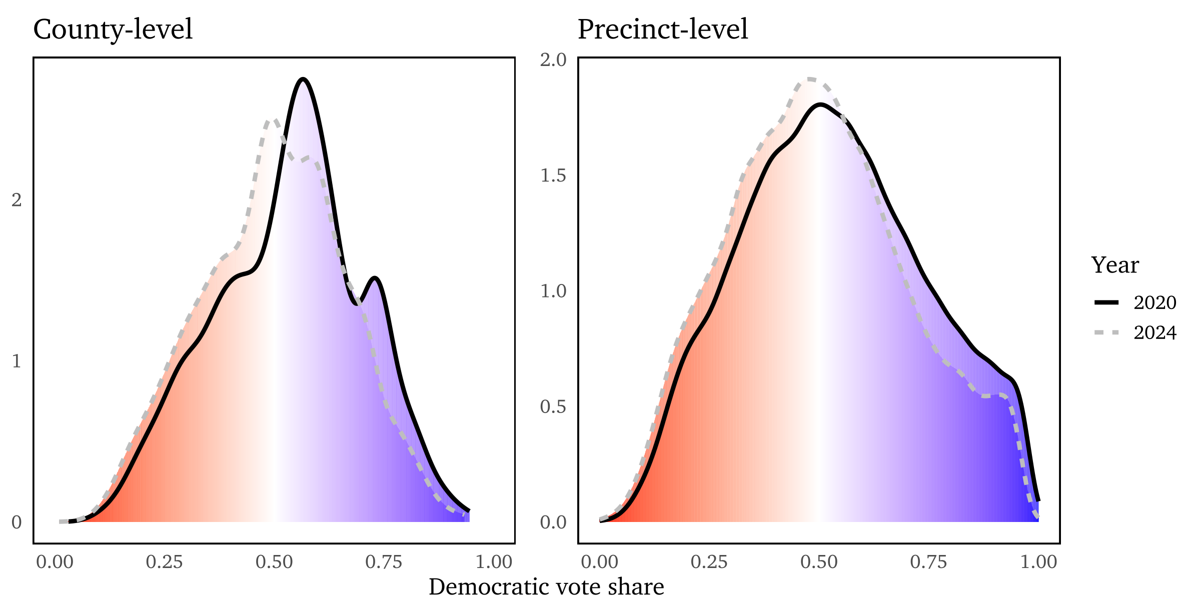

Note: Election data from the *New York Times*. The weighted standard deviation of the county-level distribution decreases from 0.167 to 0.162 (left panel). The weighted standard deviation of the precinct-level distribution decreases from 0.209 to 0.204 (right panel).

A.3 LA County Map and Cartogram

Figure A.3: Mapping Expected Kullback-Leibler Divergence in LA County (Cartogram)

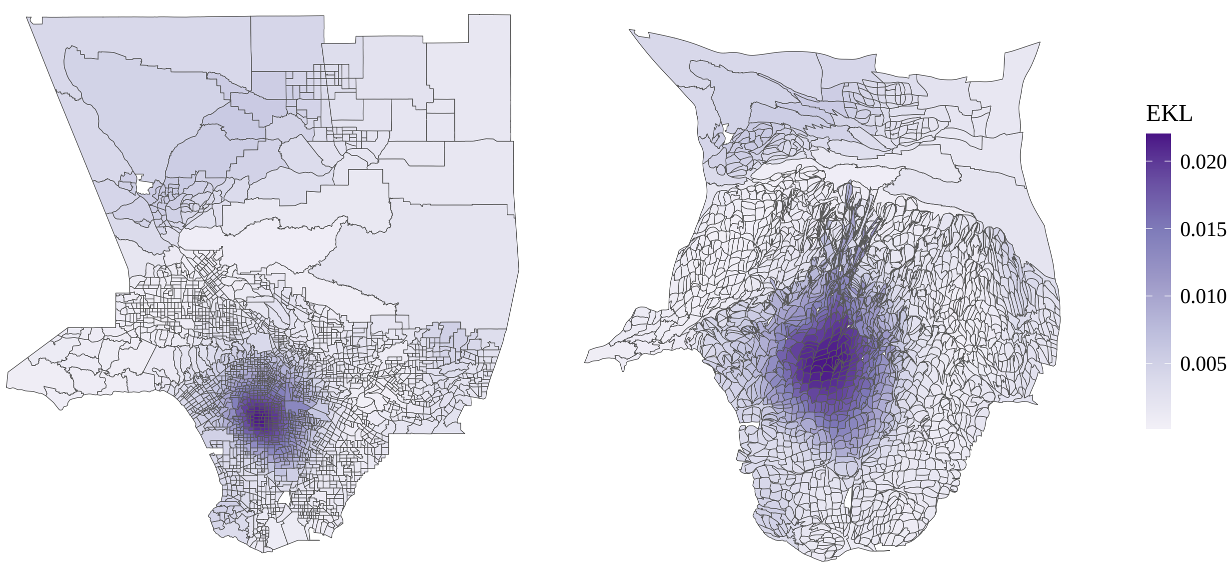

Note: Figure shows map of LA County (left) and cartogram wherein the tracts are sized proportional to their total number of registered voters. This makes tracts in the lesser populated mountainous regions of LA County smaller and the urban center larger.

A.4 Model Fitting

A.4.1 Trace Plots

Figure A.4: Orange County Fit

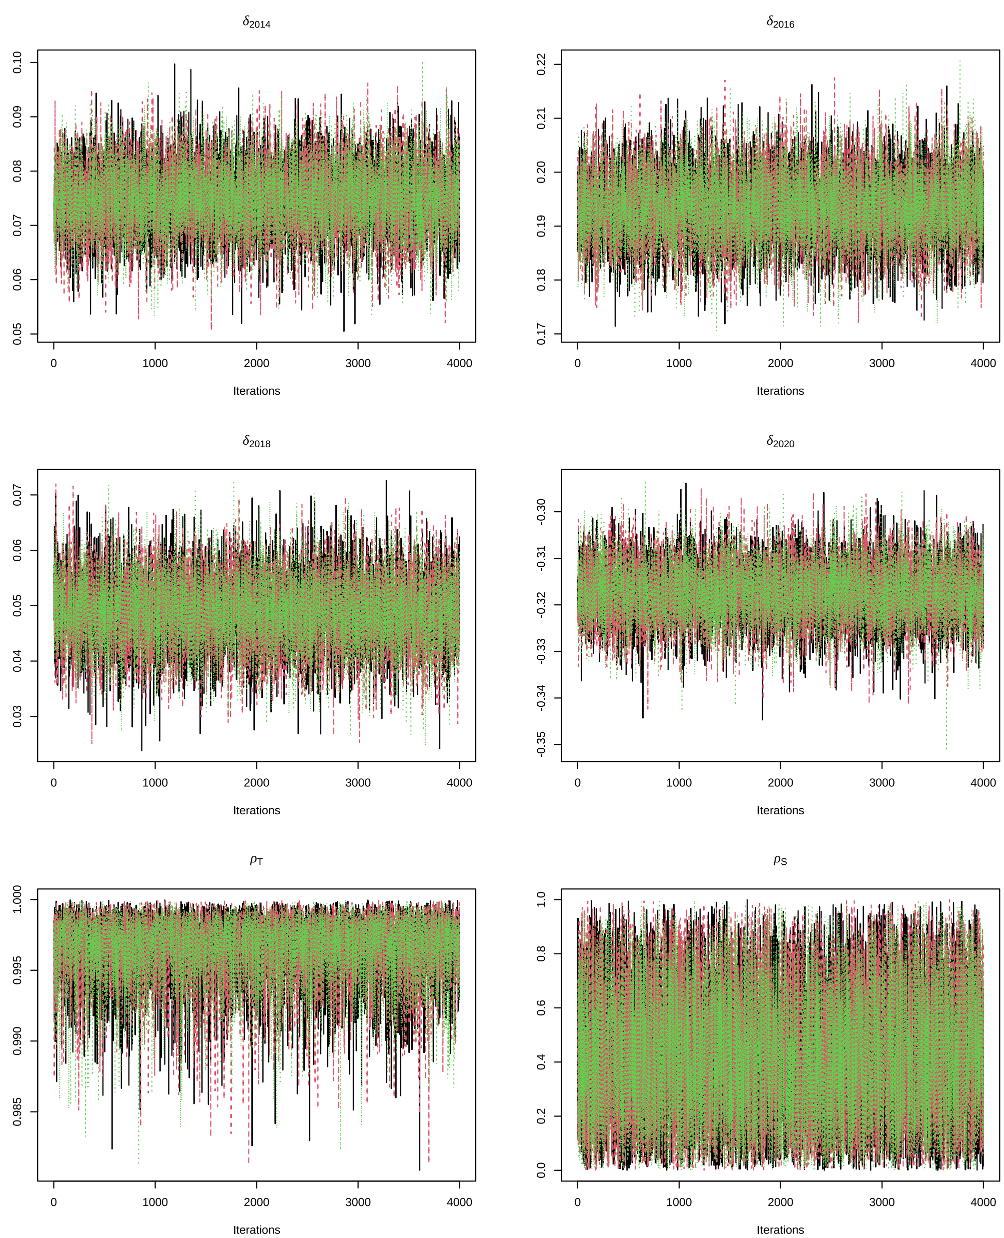

A.5 Voter Files and Self-Identification

Figure A.5: Voter File Registration-Based Partisanship Among Self-Identified Partisans According to CES (Schaffner, Ansolabehere, and Shih 2023) Survey Data

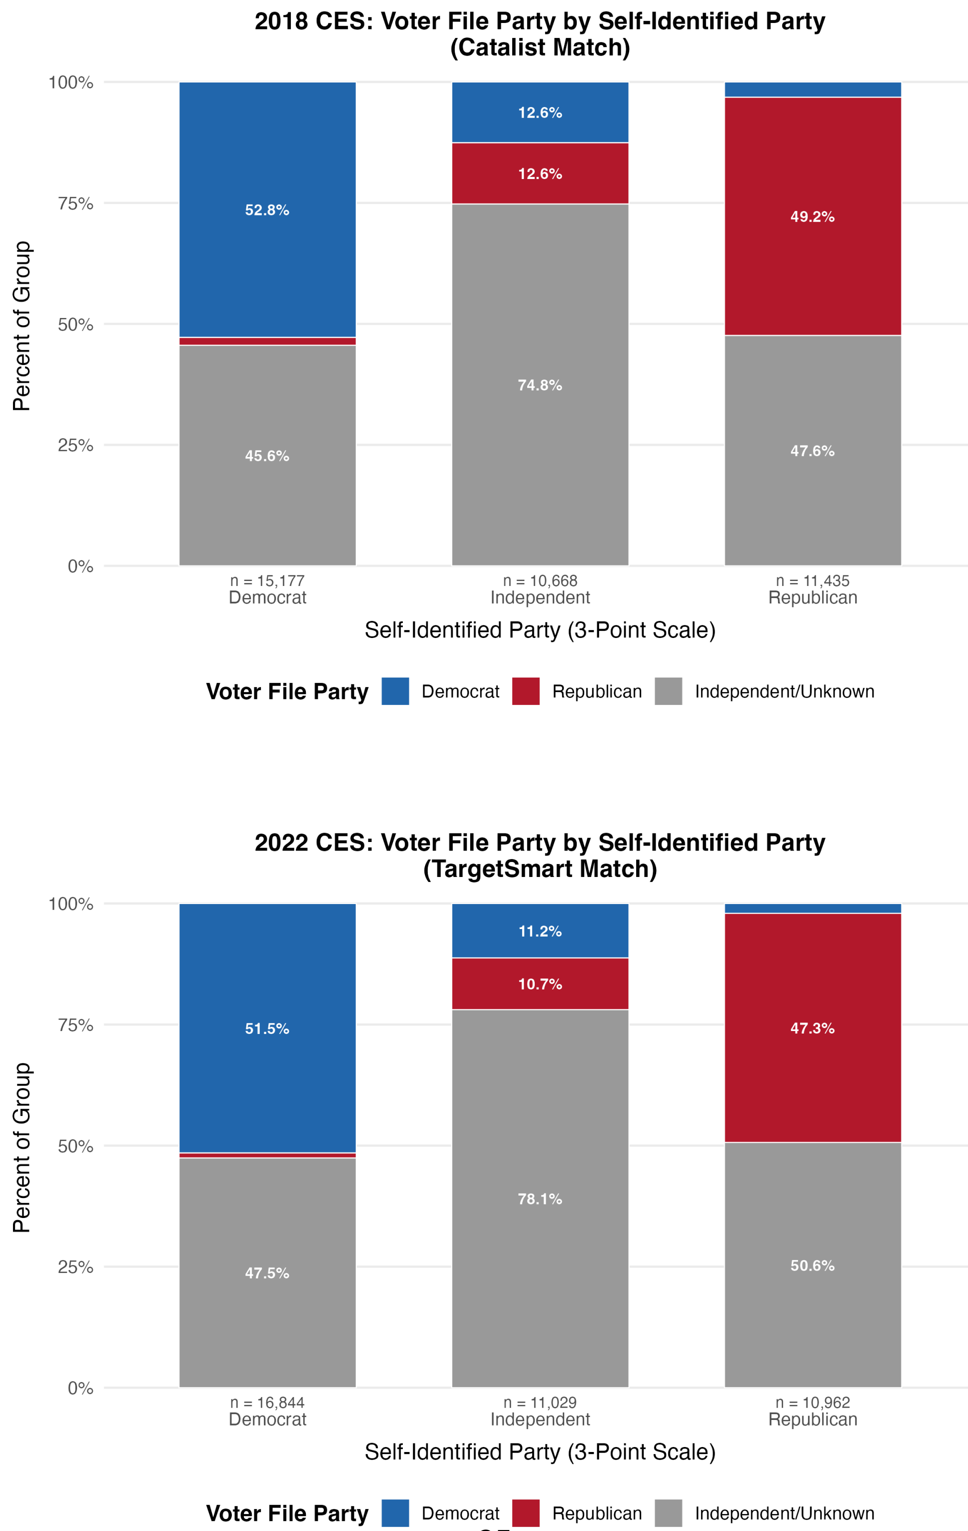

References

Abrams, Samuel J., and Morris P. Fiorina. 2012. "“The Big Sort” That Wasn't: A Skeptical Reexamination." *PS: Political Science & Politics* 45(April): 203–210.

Atske, Sara. 2019. "Political Independents: Who They Are, What They Think.".

Badger, Emily, Kevin Quealy, and Josh Katz. 2021. "A Close-Up Picture of Partisan Segregation, Among 180 Million Voters." *The New York Times* (March).

Besag, Julian, Jeremy York, and Annie Mollié. 1991. "Bayesian Image Restoration, with Two Applications in Spatial Statistics." *Annals of the Institute of Statistical Mathematics* 43(March): 1–20.

Bishop, Bill. 2009. *The Big Sort: Why the Clustering of Like-Minded American Is Tearing Us Apart*. Boston: Mariner Books.

Brown, Jacob R., and Ryan D. Enos. 2021. "The Measurement of Partisan Sorting for 180 Million Voters." *Nature Human Behaviour* 5(March): 998–1008.

Brown, Jacob R., Enrico Cantoni, Ryan Enos, Vincent Pons, and Emilie Sartre. 2025. "Sources and Extent of Rising Partisan Segregation in the U.S. – Evidence from 143 Million Voters.".

Chen, Jowei, and Jonathan Rodden. 2013. "Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures." *Quarterly Journal of Political Science* 8(June): 239–269.

Enos, Ryan D. 2017. *The Space between Us: Social Geography and Politics*. Cambridge: Cambridge University Press.

Gimpel, James G., and Iris S. Hui. 2015. "Seeking Politically Compatible Neighbors? The Role of Neighborhood Partisan Composition in Residential Sorting." *Political Geography* 48(September): 130–142.

Handcock, Mark, Julien Randon-Furling, Madalina Olteanu, and Cécile de Bezenac. 2019. "Notes on Segregation through a Multiscalar Lens.".

Ihlanfeldt, Keith, and Cynthia Fan Yang. 2024. "The Role of Neighborhood Characteristics in Explaining Political Party Residential Segregation." *Regional Science and Urban Economics* 105(March): 103992.

Kaplan, Ethan, Jörg L. Spenkuch, and Rebecca Sullivan. 2022. "Partisan Spatial Sorting in the United States: A Theoretical and Empirical Overview." *Journal of Public Economics* 211(July): 104668.

Kaysen, Ronda, and Ethan Singer. 2024. "Millions of Movers Reveal American Polarization in Action." *The New York Times* (October).

Lee, Duncan, Alastair Rushworth, and Gary Napier. 2018. "Spatio-Temporal Areal Unit Modeling in R with Conditional Autoregressive Priors Using the CARBayesST Package." *Journal of Statistical Software* 84(9): 1–39.

Leroux, Brian G., Xingye Lei, and Norman Breslow. 2000. Estimation of Disease Rates in Small Areas: A New Mixed Model for Spatial Dependence. In *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, ed. M. Elizabeth Halloran, and Donald Berry. New York, NY: Springer pp. 179–191.

Massey, Douglas S., and Nancy A. Denton. 1988. "The Dimensions of Residential Segregation." *Social Forces* 67(2): 281–315.

Mummolo, Jonathan, and Clayton Nall. 2017. "Why Partisans Do Not Sort: The Constraints on Political Segregation." *The Journal of Politics* 79(January): 45–59.

Olteanu, Madalina, Julien Randon-Furling, and William A. V. Clark. 2019. "Segregation through the Multiscalar Lens." *Proceedings of the National Academy of Sciences* 116(June): 12250–12254.

Reardon, Sean F., and Glenn Firebaugh. 2002. "Measures of Multigroup Segregation." *Sociological Methodology* 32: 33–67.

Santoro, Erik, and David E. Broockman. 2022. "The Promise and Pitfalls of Cross-Partisan Conversations for Reducing Affective Polarization: Evidence from Randomized Experiments." *Science Advances* 8(June): eabn5515.

Schaffner, Brian, Stephen Ansolabehere, and Marissa Shih. 2023. "Cooperative Election Study Common Content, 2022.".

Schroeder, Jonathan, David Van Riper, Steven Manson, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles. 2025. "National Historical Geographic Information System: Version 20.0.".

Stromberg, David. 2008. "How the Electoral College Influences Campaigns and Policy: The Probability of Being Florida." *American Economic Review* 98(June): 769–807.

Sussell, Jesse. 2013. "New Support for the Big Sort Hypothesis: An Assessment of Partisan Geographic Sorting in California, 1992–2010." *PS: Political Science & Politics* 46(October): 768–773.

Thomsen, Jens Peter Frølund, and Anna Håland Thomsen. 2023. "Intergroup Contact Reduces Affective Polarization but Not among Strong Party Identifiers." *Scandinavian Political Studies* 46(4): 241–263.

Tilley, James, and Sara B. Hobolt. 2025. "The Effect of Politically Homogenous Neighbourhoods on Affective Polarization: Evidence from Britain." *European Journal of Political Research* 64(2): 930–942.

Warner, Seth. 2025. "Affective Polarization across Time and Place in the United States.".