

SSUS WORKING PAPER SERIES

A Spatial Concordance Index of Income Segregation

Elijah Knaap[1] · Sergio J. Rey[2]

[1]University of California, Irvine, Urban Planning and Public Policy [2]San Diego State University, Geography

Spatial Science and Urban Systems · UC Irvine 2026-11-03

ABSTRACT Income segregation and income polarization describe the same phenomenon, the separation between portions of the income distribution, yet their literatures remain largely disparate. Segregation indices, led by Reardon's rank-order measures, cover the full income distribution through the composition of each neighborhood; spatial polarization indices treat space as connectivity, asking whether the groups a cut of the distribution defines form distinct clusters on the map. We show that most segregation, polarization and spatial autocorrelation measures are cross-product statistics or functionals of the same cross-product matrix, placed in one table by how they build the attribute and how they summarize the resulting graph. Following, we introduce K, a spatial income concordance index: the mean of the Perron roots of the clusters formed by neighboring units that depart from the regional composition in the same direction, weighted by each cluster's number of units, normalized by an explicit family of arrangements and corrected for chance. Across California's twenty-one metropolitan areas, the K index shows every region is most organized at a cut isolating an affluent minority. Between the 2008-2012 and 2017-2021 American Community Surveys the summary fell in sixteen of nineteen regions, by a median of 16 percent, in every band of the income distribution and least at its top, while seventeen of the nineteen came to peak in the upper half, against eleven a decade earlier: concentrated affluence became the clearly dominant pattern where regions are most organized as the spatial organization of income weakened more broadly, a restructuring that Reardon and O'Sullivan's indices register less consistently.

1 Introduction

There is much we want to know about income segregation and its change over time and space. Segregation by income grew across U.S. metropolitan areas from 1970 to 2009,

most rapidly where income inequality itself rose (Bischoff & Reardon, 2014; Reardon & Bischoff, 2011), and the concentration of income extremes among neighborhoods has proved persistent (Sampson et al., 2015). Whether the trend continued, which portions of the income distribution it separates, and how that separation is arranged on the map are open questions whose answers depend on the measures used to ask them. Indeed, sociologists spent much of the last decade testing whether the rise after 2000 was an artifact of sampling variance in the estimates used to detect it (Logan et al., 2018; Reardon et al., 2018).

Compared with the literature on racial segregation, the literature on income segregation is relatively meager (Jargowsky, 1996). Its measurement has converged on the rank-order indices of Reardon (2011), which compute a binary segregation index at every threshold of the income distribution and integrate the resulting curve into a single summary, H^R. The rank-order information theory index has become the field's canonical measure, the one on which the trend estimates above rest. It also extends to spatial measurement. Instead of treating each census tract as a self-contained neighborhood, a spatial index widens each unit's neighborhood outward to include the units around it, through a spatial weights matrix that we describe in Section 4.

The literature on polarization addresses largely the same subject. Where inequality measures the dispersion of a distribution, polarization measures its division into groups whose members resemble one another and differ from the members of other groups (Duclos et al., 2004; Esteban & Ray, 1994; Wolfson, 1997), so that the two can move in opposite directions. In recent work, Rey (2026a) extends the concept to *spatial polarization*, asking whether the groups a cut of the income distribution defines also form contiguous clusters on the map, and Rey (2026b) supplies its sampling theory.

These literatures remain largely disparate, though both are concerned with separation between portions of the income distribution. They differ in how they make the concept operational. The segregation literature measures the departure of each unit's income composition from the region's, at every threshold of the distribution, and summarizes the departures by their variance or their entropy. The polarization literature fixes the groups, usually by splitting the distribution at its median, and measures how distinct they are; its spatial extension asks whether each group's members are connected to one another on the map. Each therefore holds something the other lacks. The rank-order construction covers the full distribution, where the polarization indices evaluate one cut of it. Spatial polarization treats space as connectivity, so that separation is the division of the map into distinct clusters and isolation is a group's confinement to clusters of its own, where the spatial segregation indices average each unit's composition over a smoothed neighborhood. The two literatures even make the same methodological argument, that the curve over the distribution carries more than any scalar summary of it, without citing one another (Reardon, 2011; Wolfson, 1997).

Here we bring the two literatures together, with a bridge into spatial statistics, to combine the segregation literature's view of the full distribution with the more nuanced concept of space, separation, and isolation in spatial polarization. The bridge rests on

one observation: most segregation, polarization, and spatial autocorrelation measures are cross-product statistics or functionals of the same cross-product matrix, so a single table places them by how each builds its attribute and how each summarizes the resulting graph. One cell of that table is empty. We fill it with K, an index of spatial income concordance computed from the Perron roots of the clusters formed by neighboring units departing from the regional income composition in the same direction, and with its summary over the income distribution, K^R, the analogue of Reardon's H^R. Where the variance ratio index measures the exposure of a region's households to their own side of an income threshold, K measures how that exposure clusters: whether like tracts form contiguous enclaves, and how coherent those enclaves are.

Reardon and O'Sullivan fold clustering into evenness, on the argument that the two coincide once proximity is measured continuously (Reardon & O'Sullivan, 2004); on the discrete units that census data supply the two can diverge. The clustering matters in its own right, because public goods in the United States are funded and delivered locally: a contiguous cluster of affluent households that coincides with a municipality or a school district separates resources as well as residents (Owens et al., 2016; Trounstine, 2020). K also splits exactly into the two sides of each income threshold, so it reports which side's clustering carries a region's spatial organization, and sweeping the threshold locates where in the income distribution that organization peaks.

We apply the index to the twenty-one California metropolitan areas holding at least fifty census tracts in the 2017–2021 American Community Survey, and compare the nineteen that hold at least fifty tracts in the 2008–2012 survey as well. In the cross-section, nearly every region is most organized at a cut that isolates an affluent minority, and at that cut the affluent side carries most of the index, in keeping with the finding that affluence is more segregated than poverty (Reardon & Bischoff, 2011). Across the whole income distribution the affluent side leads as well, in fourteen of the seventeen regions whose index departs from chance, though by a smaller margin; the low-income side leads only in San Francisco, San Jose and Modesto.

Over the decade, K^R fell in sixteen of the nineteen regions, by a median of 16 percent, so by our measure spatial income segregation generally declined. The decline reached every band of the income distribution and was smallest above the seventieth percentile, while the cut at which regions are most organized moved up the distribution: seventeen of the nineteen now peak in its upper half, against eleven a decade earlier. Put differently, concentrated affluence became the clearly dominant pattern where regions are most organized, while the spatial organization of income weakened more broadly, a restructuring across the whole metropolitan area of the kind Hwang et al. (2026) describe for segregation between race–income groups. Reardon and O'Sullivan's spatial indices, computed on the same graphs, register a less consistent change, falling in only ten or eleven of the nineteen regions.

In the sections below we briefly review the literatures on spatial segregation (Section 2) and spatial polarization (Section 3), and their links to spatial statistics and spatial econometrics, showing that most of their measures belong to one generalized

family of cross-product statistics (Section 4). We then introduce K and its summary K^R (Section 5), apply them to California's metropolitan areas in 2012 and 2021 and describe how the regions changed (Section 6), and conclude with the computational problems the index leaves open and directions for future work (Section 7).

2 Income Segregation and Its Spatial Measures

Consider a region partitioned into n areal units indexed by i, each holding t_i households, with $T = \sum_i t_i$. Fix a threshold, or cut, p in the region's income distribution, let $Q(p)$ denote the income below which a share p of the region's households fall, and let s_i denote the share of unit i's households falling below $Q(p)$. The regional share is the population-weighted mean $q = \sum_i t_i s_i / T$, and the ratio s_i/q is the location quotient of the below-threshold population in unit i. It asks the question regional science has always asked of a subpopulation: is this group more concentrated here than in the region as a whole? Most segregation indices are functions of it (Massey & Denton, 1988; Reardon & Firebaugh, 2002). We work with its difference form,

$$z_i = s_i - q, \tag{1}$$

whose population-weighted mean is exactly zero, since $\sum_i t_i s_i$ and qT are both the regional count of households below $Q(p)$. The population above the threshold has location quotient $(1 - s_i)/(1 - q)$, and because the two shares sum to one, exactly one of the two quotients exceeds one wherever $s_i \neq q$. A unit's side of the cut is the group whose quotient exceeds one, and the sign of z_i records it, positive for households below the cut and negative for households above it.

The attribute already carries a segregation index before any geography enters. Its population-weighted second moment, normalized by the regional composition,

$$P = \frac{\sum_i t_i z_i^2}{T\, q(1-q)}, \tag{2}$$

is the variance ratio index. White (1986) shows it identical for a dichotomy to the correlation ratio η^2, on which Jargowsky (1996) built his measure of economic segregation, and Reardon & Firebaugh (2002) derive it as the normalized exposure index. We write it P, after that last reading, and reserve R for Reardon's rank-order variance ratio index and Reardon and O'Sullivan's spatial relative diversity index below. P reports the share of the variance in income status lying between units rather than within them, and it sits on Massey & Denton (1988)'s exposure dimension.

Evaluated at one threshold, such an index treats income as a dichotomy. Reardon (2011) generalizes it to the whole distribution: compute the binary index at every threshold, giving a curve $H(p)$, and integrate the curve against the binary entropy $E(p) = -[p \log_2 p + (1-p)\log_2(1-p)]$,

$$H^R = 2 \ln 2 \int_0^1 E(p)\, H(p)\, dp, \tag{3}$$

which weights each threshold by the entropy of the dichotomy it defines, so that cuts near the median, where the split is most even, count the most. The same construction applied to the variance ratio, with weights proportional to $p(1-p)$ in place of the entropy, gives the rank-order variance ratio index R^R, and the curves themselves show where in the distribution segregation is greatest. Measured this way, income segregation rose between 1970 and 2009, rose most where inequality grew, and separated affluent households more sharply than poor ones (Bischoff & Reardon, 2014; Reardon & Bischoff, 2011).

The indices so far treat each unit as a self-contained neighborhood, indifferent to where the units sit. Relabelling the units at random leaves Equation 2 unchanged, so a map with its high-share units gathered in one quarter of the region scores exactly what the same units scattered would. A spatial index widens the neighborhood. Reardon & O'Sullivan (2004) replace each unit's composition with that of its local environment, a proximity-weighted average of the compositions around it,

$$\tilde{s}_i = \sum_j \tilde{w}_{ij}\, s_j, \qquad \tilde{w}_{ij} = \frac{\phi_{ij}\, t_j}{\sum_k \phi_{ik}\, t_k}, \tag{4}$$

where ϕ_{ij} is a proximity function, a distance kernel or an indicator of adjacency, and compute the segregation index on the environments in place of the units. Their spatial information theory index $\tilde{H}$, spatial relative diversity index $\tilde{R}$ and spatial exposure index $\tilde{P}^*$ follow from that substitution. Sweeping the bandwidth of ϕ yields the segregation profile of Reardon et al. (2008), which reports segregation as a function of the neighborhood's radius and descends toward zero as every environment converges on the regional composition. Reardon and O'Sullivan also reduce Massey & Denton (1988)'s five dimensions to two, spatial exposure and spatial evenness, on the argument that evenness and clustering are one dimension once proximity is measured continuously. Other spatial measures of income segregation follow the same logic with different summaries, among them the spatial ordering index of Dawkins (2007), a ratio of two covariances, and the spatial counterparts of the dissimilarity and neighborhood sorting indices examined by Rey & Folch (2011).

These measures compute the variance, the entropy, or the covariance of the smoothed compositions: the difference between the typical neighborhood and the region, reported as one number for both tails of the distribution at each threshold. The polarization literature's spatial extension asks a different question: whether like units form connected clusters, and on which side of the threshold those clusters fall.

3 Polarization and Spatial Polarization

Inequality and polarization are distinct concepts. Inequality measures the dispersion of a distribution; polarization measures its division into groups, each internally similar and distant from the others. Esteban & Ray (1994) axiomatize polarization as the aggregate antagonism between groups, rising with each group's internal identification and with the alienation between groups, and Duclos et al. (2004) extend the measure to continuous distributions. Wolfson (1997) builds his polarization curve as Lorenz built his, from incomes ordered by rank and cumulated, but cumulated outward from the median instead of upward from the bottom, so that the curve measures the spread of incomes away from the middle. A progressive transfer confined to one side of the median lowers inequality and raises bipolarization, so the two can diverge. Wolfson also insists that the curve, not the scalar index summarizing it, is the object of analysis, since two distributions can share an index while polarizing at opposite ends.

Polarization between places has usually been measured on groups fixed in advance. Zhang & Kanbur (2001), for instance, take a country's geographic regions as the groups and measure polarization between them. Rey (2026a) observes that one family of measures defines groups by position in the income distribution and the other by geography, and that neither does both. His spatial polarization index builds two graphs over the same units: an attribute graph joining units in the same income class and a geographic graph joining neighbors. Their intersection joins neighbors in the same class, and its connected components are spatially coherent income clusters, endogenous in both dimensions. With k_I the number of components of the intersection and k the larger of the two graphs' own component counts,

$$S = 1 - \frac{k_I - k}{n - k}, \tag{5}$$

which reaches one when the intersection has no more components than the classes force and zero when every unit stands alone. At a median split, written S^B, it measures bipolarization in space, and it factors into contributions from the two sides of the cut. The share of geographic edges joining units of the same class gives a global homogeneity measure that decomposes exactly into local terms. Rey (2026b) derives the sampling distribution of S from percolation theory, since S is a functional of a component count and not a quadratic form.

Spatial polarization brings a concept of space as connectivity. Separation is the division of the map into distinct clusters, and isolation is a class's confinement to clusters of its own, so the measure depends on the extent and shape of each cluster, which a smoothed environment averages over. It lacks the rank-order view: it evaluates one cut, and it uses a binary attribute, so a tract barely below the median and a tract far below it count alike. The spectral segregation index of Echenique & Fryer (2007) shares both features: it measures the clusters a group's members form, using the Perron roots of the graph they induce, again on a binary attribute.

The two literatures thus hold complementary halves of one measurement problem. The segregation literature uses the full distribution and the magnitude of each unit's departure from the regional composition, and defines the neighborhood as a smoothed average. Spatial polarization treats the neighborhood as a graph and separation as the division of that graph into clusters, at one cut of a binarized distribution. The next section places both in one frame, borrowed from spatial statistics.

4 A Common Frame: Cross-Product Statistics

Spatial statistics supplies the frame. Mantel (1967) and L. J. Hubert et al. (1981) write a general statistic of association between two matrices defined over the same units,

$$\Gamma = \sum_i \sum_j w_{ij}\, y_{ij}, \tag{6}$$

where the spatial weights matrix $W = [w_{ij}]$ records which pairs of units are near one another and $Y = [y_{ij}]$ records how alike each pair is in some attribute. Moran's I, Geary's c, the join-count statistics and the Getis–Ord statistics are instances, differing only in how Y is built from the attribute (Cliff & Ord, 1973; Getis, 1991; Getis & Ord, 1992). For an attribute x with mean $\bar{x}$, Moran's I takes $y_{ij} = (x_i - \bar{x})(x_j - \bar{x})$,

$$I = \frac{n}{S_0} \frac{\sum_i \sum_j w_{ij}(x_i - \bar{x})(x_j - \bar{x})}{\sum_i (x_i - \bar{x})^2}, \qquad S_0 = \sum_i \sum_j w_{ij}, \tag{7}$$

while the join counts take an indicator that i and j fall in the same class. Each statistic also has a local form, the contribution of unit i to the sum, which Anselin (1995) and Ord & Getis (1995) develop as local indicators of spatial association.

The segregation indices of Section 2 are cross-product statistics in the same sense. Write $\tilde{W} = [\tilde{w}_{ij}]$ for the population-weighted proximity matrix of Equation 4. Because each row of $\tilde{W}$ sums to one, the departure of unit i's environment from the region is $\tilde{s}_i - q = (\tilde{W}z)_i$, and the three indices become

$$P = \frac{\sum_i t_i\, z_i^2}{T\, q(1-q)}, \qquad \tilde{P} = \frac{\sum_i t_i\, z_i (\tilde{W}z)_i}{T\, q(1-q)}, \qquad \tilde{R} = \frac{\sum_i t_i\, (\tilde{W}z)_i^2}{T\, q(1-q)}, \tag{8}$$

where $\tilde{P} = (\tilde{P}^* - q)/(1-q)$ is Reardon and O'Sullivan's spatial exposure index of the below-threshold group to itself, its isolation, normalized as the variance ratio normalizes its aspatial counterpart. The last two equalities hold up to a term proportional to $\sum_i t_i (\tilde{W}z)_i$, which vanishes when the smoothing preserves the regional composition. The spatial information theory index is proportional to $\tilde{R}$ to second order in the departures, $\tilde{H} \approx \tilde{R}/2E(q)$ with the entropy taken in natural units, so the factor changes with the threshold. The spatial segregation indices therefore take two forms. The exposure index is the cross-product of a unit's own departure with its environment's, a population-weighted analogue of the numerator of Moran's I on $\tilde{W}$. The evenness indices are the squared departure of the environment, a quadratic form in $\tilde{W}'\,\mathrm{diag}(t)\,\tilde{W}$, which joins units

that share a neighbor. Put differently, exposure is a covariance between a unit and its neighborhood, and evenness is the variance of the neighborhood itself. P is their common origin, the case $\tilde{W} = I$ in which every unit is its own neighborhood.

The same two forms separate the two workhorse models of spatial econometrics. A conditional autoregressive (CAR) field on a symmetric graph A with degree matrix D has precision $D - \rho A$, so the cross-product $x'Ax$ is the interaction term of its log density (Besag, 1974); a simultaneous autoregressive (SAR) field on W has precision $(I - \rho W)'(I - \rho W)$, which adds the two-step term $W'W$. Exposure-type measures thus correspond to the interaction term a CAR model specifies, and evenness-type measures to the term a SAR model adds to it. We return to the CAR field in Section 5, where it defines the normalizer of the index.

The polarization measures use a binary attribute. Setting $y_{ij} = 1$ when units i and j fall on the same side of the threshold, and zero otherwise, turns Equation 6 into a join count; Rey's homogeneity measure is the share of geographic edges that count includes. The graph-based indices summarize the same matrix, $B = W \odot Y$, the geographic graph masked by the attribute graph, by a different functional. S counts the connected components of B, which is the number of zero eigenvalues of its Laplacian, and the spectral segregation index takes the Perron root, the largest eigenvalue, of each component (Echenique & Fryer, 2007). A sum of B's entries counts adjacent like pairs; an eigenvalue measures the assembly of those pairs into clusters.

Table 1: Segregation, spatial autocorrelation and polarization measures placed by how the attribute matrix Y is built and which functional summarizes $W \odot Y$, each up to its normalization. Magnitude means $y_{ij} = z_i z_j$; sign only means an indicator that i and j fall on the same side of the threshold. The spatial segregation indices carry population weights, so their $W'W$ is $\tilde{W}'\mathrm{diag}(t)\tilde{W}$, and $\tilde{H}$ enters to second order, as a multiple of $\tilde{R}$. K is computed from the truncated, household-weighted form of Equation 9.

	Attribute keeps magnitude	Attribute keeps sign only
Sum of $W \odot Y$	P (W the identity); Moran's I and $\tilde{P}$ (W); $\tilde{R}$ and $\tilde{H}$ ($W'W$)	join counts; Rey's homogeneity
Eigenvalue of $W \odot Y$	*(empty; filled by K)*	Rey's S; spectral segregation index

Two independent choices thus organize the measures (Table 1): whether the attribute matrix keeps the magnitude of each unit's departure from the regional composition or only its sign, and whether the masked graph is summarized by its sum or by an eigenvalue. A third choice, the threshold p that builds the attribute, runs across both, since the rank-order construction sweeps it and the polarization indices fix it. One cell of the table is empty: an eigenvalue of a graph weighted by the magnitudes of the departures. An index in that cell would measure clusters, as the spatial polarization indices do, while keeping the magnitudes and, swept across thresholds, the

full distribution that the segregation indices cover. Procedures that detect clusters with local statistics, such as significant high–high and low–low local Moran clusters (Anselin, 1995) or local Getis–Ord hot spots (Ord & Getis, 1995), also find concordant clusters, but one location at a time and against a significance threshold. They are local terms of Equation 6, in the sum row, and they locate clusters without summarizing a region's concordance on one scale. The next section constructs the index the empty cell calls for.

5 The Spatial Concordance Index

The construction runs in five steps, which the subsections below take in turn. At a cut in the income distribution, each unit falls on the side of the cut whose location quotient exceeds one, and neighboring units on the same side form clusters. Each cluster contributes its Perron root, which combines the cluster's departures with the ties among them, and λ averages the roots, weighting each cluster by its number of units. Because λ scales with the departures, we normalize it at every cut by the largest value an explicit family of arrangements attains and correct it for the concordance random arrangements of the same units produce. The normalized index splits exactly by side of the cut, and sweeping the cut traces the curve $K(p)$, whose peak locates the cut at which a region's income geography is most coherently organized and whose entropy-weighted integral, K^R, summarizes the whole distribution.

5.1 Concordant Clusters

Let A be a symmetric spatial adjacency or weights matrix with degrees $D = \operatorname{diag}(A\mathbf{1})$, and write $W = D^{-1/2} A D^{-1/2}$ for its row-standardized form written symmetrically, which has the same spectrum as $D^{-1}A$; for K this W takes the place of the generic weights of Section 4. Standardizing keeps the centrality of a unit in the graph from entering the index as concentration, and the symmetric form keeps the eigenvalues real. We weight each unit's departure by its household count, $z_i^{(t)} = z_i \sqrt{t_i}$, so that a tract with few households cannot dominate on the magnitude of its share alone, and form

$$M = W \odot \left[\max\left(0,\ z_i^{(t)} z_j^{(t)}\right) \right]_{ij}. \tag{9}$$

Each entry of M is the product of a tie's strength and the two departures it joins, kept when the two units depart from the regional composition in the same direction and set to zero when they depart in opposite directions. At $W = I$ every component of M is a single unit, and the λ of Equation 11 reduces to $\sum_i t_i z_i^2 / n$, the numerator of P over n, so λ extends the variance ratio index into the eigenvalue row of Table 1. The normalized index is undefined there, since a graph without edges gives every rearrangement the same λ.

The truncation fixes the geometry of M. No edge of M joins a unit that over-represents households below the cut to one that over-represents households above it, so M sep-

arates into components, each a maximal set of connected units on one side of the threshold. On a connected graph a component can end only where the side changes, so each component is a contiguous cluster of like units, an enclave in Massey & Denton (1988)'s sense. P measures exposure at the scale of the units; K measures how that exposure clusters. K is thus an exposure index in its attribute, which places it with the exposure measures on Reardon and O'Sullivan's criteria (Table 3), and a clustering index in what it measures. On a fragmented graph, by contrast, a component can also end at a gap in the graph, and components stop meaning clusters; we therefore compute K on contiguity graphs.

5.2 Perron Roots

The Perron root of a component, its largest eigenvalue, measures how coherently the component holds together. With M_k the block of M on component C_k and v_k its Perron vector, normalized to unit length, the root is the quadratic form

$$\lambda_k = v_k' M_k v_k = \sum_{i,j \in C_k} v_{ki}\, v_{kj}\, w_{ij}\, z_i^{(t)} z_j^{(t)}, \tag{10}$$

a sum over the component's ties of the product of the two departures each tie joins, weighted by the tie's strength and by the Perron vector at both ends. A unit therefore counts only through its ties to like neighbors, and the root grows with the magnitudes of the departures and with the density of the ties among them, so a compact cluster of large departures scores higher than a chain of the same units. The Perron root of M itself reports only the densest component, since the spectrum of a block-diagonal matrix is the union of its blocks' spectra: on the Los Angeles queen contiguity graph, between four and twenty-two of 3,080 tracts carry ninety percent of its Perron vector's mass, depending on the threshold. Following Echenique & Fryer (2007), we instead take the size-weighted mean of the components' roots,

$$\lambda = \frac{1}{n} \sum_k |C_k|\, \lambda_k, \tag{11}$$

where C_k are the components of M, $|C_k|$ the number of units in each, and λ_k their Perron roots, so that every concordant cluster contributes in proportion to the number of units it spans. Households enter through the departures $z_i^{(t)}$, not through $|C_k|$.

The root also grows with a cluster's extent, so K measures the spatial scale of segregation from within. A block of equal departures a has root $a^2 \rho(W_C)$, where $\rho(W_C)$ is the spectral radius of the graph restricted to the block. A unit on the block's edge loses the weight of its ties to the other side, so $\rho(W_C)$ is small for a block of two or three units and approaches one as the block's interior outgrows its boundary. On the Los Angeles queen graph its median is 0.17 to 0.20 for components of two or three tracts, 0.34 to 0.47 for four to ten, and 0.89 to 0.98 for components of more than fifty, at each of three cuts. Equal departures therefore count for more when they extend over a larger contiguous area, and the scale at which like units cohere is set by the extent of the

clusters, not by a bandwidth chosen in advance. Size thus enters λ twice, inside each root through the extent of the cluster and outside it through the weight $|C_k|$.

Measuring clusters by their Perron roots gives the index the concept of space that spatial polarization introduced. A cluster enters whole, through its extent and its internal coherence, wherever its contiguity carries it, so K treats a wedge, a corridor or a sector as one object, where a window-based index averages over whatever rings each unit's environment happens to cover. On synthetic geographies holding one fixed set of unit compositions in different shapes (Appendix A), wedges, rings and bands give K between 0.69 and 1.26 times the value the same compositions give as a compact blob, and Reardon and O'Sullivan's $\tilde{R}$ on one ring of neighbors between 0.55 and 1.02. The gap is widest on a band about two units wide, which keeps K at 0.99 times the blob's value where $\tilde{R}$ falls to 0.81 on one ring and 0.64 on two. Scattered as isolated units, the same compositions take K below chance, to about -0.09 of the blob's value, while $\tilde{R}$ on one ring still reports 0.20 to 0.30 of it, because each unit's environment includes the unit's own departure.

5.3 Normalization and Chance

λ is homogeneous of degree two in the departures, so comparing it across regions or thresholds requires a normalizer. The segregation literature divides by the maximum attainable given the marginals; here the natural maximum is the largest λ that any rearrangement of the same departures attains on the same graph. That maximum has no formula: on a regular graph with a two-valued attribute it decides whether the graph contains a clique of a given size, so computing it is NP-hard (Appendix B). We therefore define $\lambda^{\max}$ as the maximum over an explicit family of arrangements, in the spirit of normalizing a spatial statistic by its attainable extremes (de Jong et al., 1984). The family comes from the CAR field of Section 4 at its intrinsic limit, Besag's ICAR model (Besag & Kooperberg, 1995). Conditioning that field on a high value at one unit and a low value at another gives the smoothest surface consistent with two opposed observations, and sorting the departures along that surface places the largest positive departures around one pole and the largest negative departures around the other. $\lambda^{\max}$ is the largest λ over such sortings, for a fixed family of opposed pole pairs (Appendix B): the concordance the departures would attain if a spatial process on the same graph arranged them around two opposed observations. It is a deterministic function of the graph and the departures, and we compute it exactly. Where the true rearrangement maximum can be found by enumeration, the family comes within 1 percent of it in most cases and within 4 percent in all (Appendix B).

The ratio $\lambda/\lambda^{\max}$ still has a floor set by the graph, because a random arrangement of the departures also forms concordant clusters. On California's metropolitan areas that floor runs from 0.04 to 0.32 and moves more than threefold across a single sweep of the income distribution. We correct for it as the segregation literature measures departure from the segregation that random allocation of the same marginals would produce (Carrington & Troske, 1997), the correction that Cohen's κ and the adjusted

Rand index make for agreement between raters and between partitions (Cohen, 1960; L. Hubert & Arabie, 1985):

$$K = \frac{\lambda - \mathbb{E}[\lambda]}{\max(\lambda^{\max}, \lambda) - \mathbb{E}[\lambda]}, \tag{12}$$

where $\mathbb{E}[\lambda]$ is the mean of λ over random rearrangements of the units' (composition, household count) pairs. Under that null each side of the threshold occupies a uniformly random set of units, so the component structure of M follows site percolation on the graph, and we approximate $\mathbb{E}[\lambda]$ from the graph without permuting the attribute (Appendix C). K thereby takes a chance-anchored form of Reardon and O'Sullivan's scale for a spatial segregation index (Reardon & O'Sullivan, 2004), with chance where their criterion places the evenly mixed map: 1 when the map is at least as concordant as the family's best arrangement, 0 at chance, and negative under hyperintegration, a map less concordant than a random one. As with Moran's I, negative values are possible but uncommon in observed maps. The ceiling $\max(\lambda^{\max}, \lambda)$ guards against a map more concordant than every member of the family, and it binds in none of the cells we report, where the observed λ reaches at most 82 percent of $\lambda^{\max}$.

5.4 The Two Sides of the Threshold

Because no edge of M joins units on opposite sides of the threshold, K splits exactly,

$$K = K_{\text{low}} + K_{\text{hi}}, \qquad K_s = \frac{\lambda_s - \mathbb{E}[\lambda_s]}{\max(\lambda^{\max}, \lambda) - \mathbb{E}[\lambda]}, \quad s \in \{\text{low}, \text{hi}\}, \tag{13}$$

where λ_s sums $|C_k|\lambda_k/n$ over the components on side s, so that $\lambda = \lambda_{\text{low}} + \lambda_{\text{hi}}$. Each side is corrected against its own chance term over the shared denominator, so the two sides are measured on one scale and either can sit at or below chance. K_{low} is concordance among units that over-represent households below the cut, which we call the low-income side, and K_{hi} concordance among units that over-represent households above it, the affluent side. Both terms are concordance; the subscripts name a side, not a sign. K_{hi} thus measures the isolation of the affluent side in clusters of its own, K_{low} that of the low-income side, and K the separation of the map into such clusters. At a high cut the low-income side holds everyone outside an affluent minority, and at a low cut the affluent side holds everyone outside a poor one, so the labels name the group a unit over-represents and not an income class.

5.5 The Income Distribution

Sweeping the threshold across the income distribution gives the curve $K(p)$. Following Reardon's rank-order construction, we summarize it by its integral against the binary entropy over the swept thresholds, taken by the trapezoid rule,

$$K^R = \frac{\sum_b E(p_b)\,\Delta_b\,K(p_b)}{\sum_b E(p_b)\,\Delta_b}, \tag{14}$$

where Δ_b is the trapezoid width of cut b: half the distance between its two neighbors, or half the gap to its one neighbor at either end of the sweep. Equation 14 is the counterpart of Equation 3 over the range the sweep covers, which Reardon (2011) estimates instead by fitting a polynomial to $H(p)$ and integrating it over $(0, 1)$. The widths matter because the cuts are bracket edges, unevenly spaced in p; without them the mean would weight the thresholds by their number where the integral weights them by their spacing. The side indices K^R_{low} and K^R_{hi} sum to K^R in the same way. The curve also locates a peak, the threshold at which a region's income geography is most coherently organized. Because $K(p)$ is computed from every cluster at a cut, the peak is the cut at which the map as a whole is most concordantly organized, not the cut that produces the single most coherent cluster, and the split of Equation 13 then says whose clusters carry it there. The peak is a different quantity from the threshold at which P peaks. K divides the size of the departures out, so a map whose departures are proportional from one cut to the next gives a flat $K(p)$ wherever $P(p)$ peaks, and the two answer different questions: P where the distribution is most divided between units, K where that division is most coherently arranged in space.

Figure 1 shows what the index measures on four stylized geographies. With the affluent class concentrated in three compact enclaves, K^R is 0.27 and the affluent side carries 0.92 of it; with the low-income class concentrated in three pockets, K^R is 0.27 and the affluent side carries 0.03. With both, K^R is 0.37, the two sides contribute nearly equally over the sweep, and they cross near the 45th percentile. Scattering the enclave units across the map leaves every unit's composition, and so $P(p)$, unchanged, and takes K^R to −0.02. The side split names the clustered tail, but the peak alone does not: concentrated poverty peaks at the 70th percentile, in the upper half of the distribution, as concentrated affluence does. Interpreting a region's peak therefore requires both, the cut at which it falls and the side that carries K there.

5.6 Properties and Inference

Against Reardon and O'Sullivan's eight criteria for a spatial segregation index (Table 3), K satisfies scale interpretability in its chance-anchored form, through the correction and the ceiling, along with population density invariance and, to a vanishing residual, location equivalence. It fails composition invariance, as its aspatial parent does; boundary independence, as any measure computed on areal units does; and the exchange criterion, which none of the spatial indices Reardon and O'Sullivan describe meets in general. In place of additive spatial decomposability, which no measure built on an eigenvalue satisfies, it offers the exact split by side, and additive grouping decomposability does not apply to an index of two groups.

We conduct inference on K^R for the region as a whole. Each of 999 random rearrangements of a region's units, carried across every threshold of the sweep, yields a value of K^R under the null; $\lambda^{\max}$ and $\mathbb{E}[\lambda]$ are invariant to the rearrangement, so they are fixed, and the ceiling compares each rearrangement's λ with the same $\lambda^{\max}$. With $n_{\geq}$ and $n_{\leq}$ the numbers of rearrangements at or above and at or below the observed

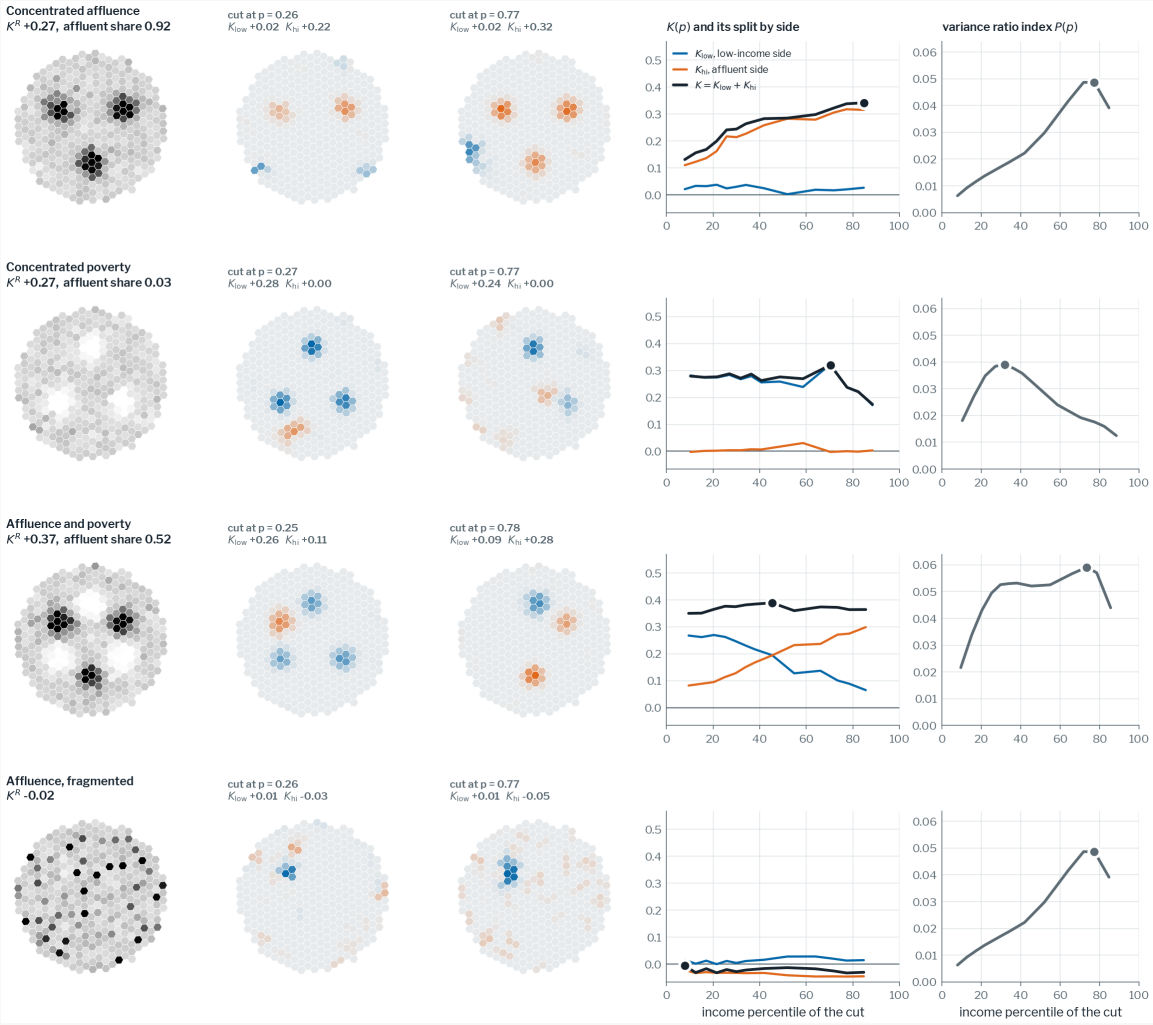


Figure 1: What K measures, on four synthetic income geographies of 360 hexagonal units. Column 1: median household income by unit (light is low), headed by the scenario's K^R and the affluent side's share of it. Columns 2 and 3: the exact split of λ among units at a low and a high cut, blue for the low-income side and orange for the affluent side, each unit shaded by its share ℓ_i of λ (Appendix E), headed by K_{low} and K_{hi} at that cut. Column 4: $K(p)$ and its two sides across the income distribution, the circle marking the peak. Column 5: the variance ratio index $P(p)$, the circle marking its peak. The fragmented scenario holds the compositions of the first scenario with its enclave units scattered, so $P(p)$ is identical.

value, the two-sided pseudo p-value is $\min\{1, 2(\min(n_{\geq}, n_{\leq}) + 1)/1000\}$. The same Perron structure divides λ exactly among units, giving a local index, but that local index concentrates on a few tracts in random rearrangements as in observed maps, so we report it only in Appendix E.

6 Income Segregation in California

6.1 Data

We compute K on California census tracts from the 2017–2021 five-year American Community Survey, table B19001, which counts each tract's households in sixteen income brackets. Of the 9,043 tracts carrying households, 8,620 fall in the twenty-one metropolitan statistical areas holding at least fifty tracts, which we call regions. The income threshold sweeps the bracket edges lying between the fifth and ninety-fifth percentiles of each region's own distribution, thirteen to fifteen cuts per region, and Equation 14 integrates $K(p)$ over those cuts. The bracket edges crowd below the median, so nine to eleven of each region's cuts fall in the lower half, and the widths in Equation 14 keep that crowding from tilting the summary. On households with known incomes binned to these brackets, the trapezoid recovers the entropy-weighted integral of the information theory index to within 3 percent; a mean over the cuts without widths misses it by up to a third, and a polynomial extrapolated past the top bracket by up to 15 percent. We compute the rank-order versions of Reardon and O'Sullivan's indices the same way. The graph is queen contiguity. Los Angeles is the one region whose contiguity graph is disconnected, and we join Santa Catalina and San Clemente islands to the mainland by a single edge each, between nearest tracts. For the comparison over time we repeat the analysis on the 2008–2012 survey, on its own tract boundaries and contiguity graphs, for the nineteen regions holding at least fifty tracts in both years. Because K, like every index computed on areal units, depends on the partition into units (Appendix D), part of any change between the years can reflect redrawn tracts. We report three populations and label each count by its population: the twenty-one regions of the 2021 cross-section; the seventeen of those whose K^R departs from chance; and the nineteen regions holding at least fifty tracts in both surveys, of which sixteen depart from chance in both.

6.2 Levels and Agreement with the Rank-Order Indices

K^R departs from its permutation distribution at the 0.05 level, two-sided, in seventeen of the twenty-one regions, each time above it (Table 2). San Luis Obispo and Merced fall just short, at pseudo p-values of 0.09 and 0.11, and Redding and Santa Rosa are indistinguishable from a random arrangement of their tracts; we exclude these four from the counts that follow except where noted. Stockton's income geography is the most concordant, at $K^R = 0.53$, followed by Vallejo, Riverside, Fresno and Salinas.

Table 2: Rank-order indices for California's twenty-one metropolitan areas, ACS 2017–2021, queen contiguity, thresholds between the fifth and ninety-fifth percentiles. K^R is Equation 14, with its two-sided pseudo p-value from 999 whole-map rearrangements (smallest attainable 0.002). $\tilde{H}^R$ and $\tilde{R}^R$ are Reardon and O'Sullivan's spatial information theory and relative diversity indices on the same graph, thresholds and weights, each tract's environment its queen neighborhood. The peak percentile is the cut at which $K(p)$ is largest; in eight regions the runner-up cut lies within 0.02 of it. The affluent share is K^R_{hi}/K^R, the fraction of K^R carried by tracts that over-represent households above the cut, averaged over the sweep.

Region	Tracts	K^R	Pseudo p-value	$\tilde{H}^R$	$\tilde{R}^R$	Peak percentile	Affluent share of K^R
Stockton	173	0.529	0.002	0.057	0.068	74	0.69
Vallejo	98	0.289	0.002	0.033	0.040	77	0.65
Riverside	978	0.285	0.002	0.057	0.069	82	0.56
Fresno	225	0.244	0.002	0.075	0.087	86	0.68
Salinas	100	0.233	0.002	0.030	0.036	79	0.63
Los Angeles	3,080	0.227	0.002	0.060	0.072	76	0.60
San Diego	730	0.224	0.002	0.053	0.063	75	0.66
Sacramento	561	0.198	0.002	0.047	0.057	60	0.51
Santa Cruz	69	0.194	0.006	0.026	0.030	15	0.54
San Francisco	1,091	0.176	0.002	0.055	0.067	72	0.41
San Jose	420	0.176	0.004	0.037	0.047	67	0.37
Bakersfield	232	0.175	0.002	0.077	0.088	60	0.64
Oxnard	188	0.166	0.004	0.028	0.038	84	0.50
Chico	54	0.162	0.014	0.010	0.018	59	0.67
Visalia	102	0.150	0.002	0.026	0.033	95	1.00
Modesto	112	0.119	0.012	0.019	0.023	93	0.35
Santa Maria	105	0.113	0.004	0.052	0.062	85	0.65
San Luis Obispo	68	0.062	0.090	0.018	0.021	11	0.57
Merced	63	0.059	0.110	0.014	0.019	95	0.27
Redding	50	0.040	0.530	0.006	0.009	94	0.46
Santa Rosa	121	0.034	0.228	0.013	0.015	74	0.54

K^R tracks the established rank-order indices without reproducing them. Computed on the same graph and thresholds, Reardon and O'Sullivan's spatial information theory and relative diversity indices order the twenty-one regions almost identically (Spearman $\rho = 0.996$), and K^R tracks both at $\rho = 0.70$. The aspatial rank-order variance ratio R^R, K's parent swept across the distribution, agrees closely with $\tilde{R}^R$ ($\rho = 0.94$), so the smoothed indices largely reproduce the aspatial ordering, while K^R agrees with R^R at $\rho = 0.62$. The departures are concentrated. Vallejo, Salinas, Santa Cruz and Chico rank well above their position on $\tilde{H}^R$, and Bakersfield and Santa Maria well below. Peak

location and the split by side are unrelated to these departures, and region size only weakly related (Spearman $\rho = -0.40$), so what drives them remains open.

6.3 Where Each Region Is Most Organized

Income geography is most organized at cuts that isolate an affluent minority (Figure 3 a). No region peaks at the median, where Wolfson's index and Rey's S^B fix the cut by assumption (Rey, 2026a; Wolfson, 1997). Sixteen of the seventeen regions with significant K^R peak in the upper half of their own distributions, from Chico, Sacramento and Bakersfield near the sixtieth percentile to Visalia at the ninety-fifth; Santa Cruz alone peaks low, at the fifteenth. Six of the sixteen, Visalia, Modesto, Oxnard, Santa Maria, San Francisco and San Jose, peak at the highest cut their sweep reaches, with $K(p)$ still rising there. Table B19001's top bracket begins at $200,000, which falls lowest in the richest regions, at the 67th percentile in San Jose and the 73rd in San Francisco. In these six regions the peak percentile and the peak value of K are therefore lower bounds, and brackets finer at the top would likely raise both, placing the most organized cut among a still smaller affluent minority. The curve need not rise without limit, since a thinning affluent set breaks into smaller blocks with smaller roots; in the other ten regions $K(p)$ peaks inside the sweep and falls toward its top, as in Los Angeles (Figure 2). The peak percentile is a coarse reading, since in eight of the twenty-one regions the runner-up cut lies within 0.02 of the peak, but the half is rarely in doubt: in twelve of the sixteen upper-half regions the best upper cut exceeds the best lower cut by at least 0.05. As Figure 1 shows, an upper-half peak alone does not identify concentrated affluence. Which side carries K at the peak does, and among the seventeen regions with significant K^R the affluent side carries a median 0.82 of K there.

Over the whole distribution the affluent side carries more as well, though by a smaller margin. Because $K = K_{\text{low}} + K_{\text{hi}}$ at every threshold, K^R splits into K^R_{low} and K^R_{hi}, and the affluent side's share K^R_{hi}/K^R does not depend on where the curve peaks. That share has a median of 0.57 across the twenty-one regions and is unrelated to peak location (Spearman $\rho = -0.02$; Figure 3 b). Among the seventeen regions with significant K^R, the affluent side leads in fourteen, most strongly in Visalia (1.00 of K^R), Stockton (0.69) and Fresno (0.68), and the low-income side in three, San Francisco, San Jose and Modesto, where the affluent side carries 0.35 to 0.41; in Oxnard and Sacramento the two sides nearly balance. The two readings answer different questions. Where K peaks, and which side carries it there, says at which cut a region is most organized and around whom; the split over the sweep says which side's clustering carries the region's organization across all cuts. Concentrated affluence describes both, the peak far more sharply: among the seventeen, the affluent side carries a median 0.82 of K at the peak and 0.63 of K^R over the sweep.

Figure 2 shows what the index measures in one region, Los Angeles, at four cuts: the bracket edges nearest its 25th and 50th percentiles, $40,000 and $75,000, the cut where its $K(p)$ peaks, $150,000, the 76th percentile, and $200,000, the 86th percentile and the highest cut the brackets allow. At the first three cuts each side of the threshold

forms one dominant connected component and many small ones. On the affluent side a single component, reaching from the coast through the foothills and across Orange County, carries 95 to 100 percent of that side's share of λ; on the low-income side one component centered on central and South Los Angeles carries 76 to 84 percent, and between 59 and 97 smaller components, among them separate clusters in the Antelope Valley, the San Fernando Valley and central Orange County, carry the rest. Within each component the Perron vector concentrates on a core a few tracts wide, so each component's root, weighted by its number of tracts, enters λ through its own core. Through $150,000 the affluent component's core lies in the master-planned communities of southern Orange County, Ladera Ranch, Rancho Mission Viejo and Coto de Caza, where eight adjoining tracts carry 81 to 91 percent of the component's λ. Those tracts are large and uniformly affluent and border little else, so they form the region's most coherent affluent block; older and richer enclaves border more mixed tracts. The core stays in place when the household weighting is removed.

The cuts differ in the coherence of each side's clusters against chance. At $40,000 the low-income side's clusters are the more coherent, with K_{low} at 0.10 and K_{hi} at 0.07; at $150,000 the affluent side carries the index, with K_{hi} at 0.27 and K_{low} at 0.04. As the cut rises, the affluent side's clusters grow more distinct and the low-income side's less so, and K climbs from 0.17 at $40,000 to its peak of 0.31 at $150,000. At $200,000 the affluent component splits in three: a Westside component centered on Pacific Palisades and Brentwood, which now carries half of that side's share of λ, southern Orange County, and the South Bay beach cities. K holds at 0.31, so the affluent side is about as coherent in three blocks as in one.

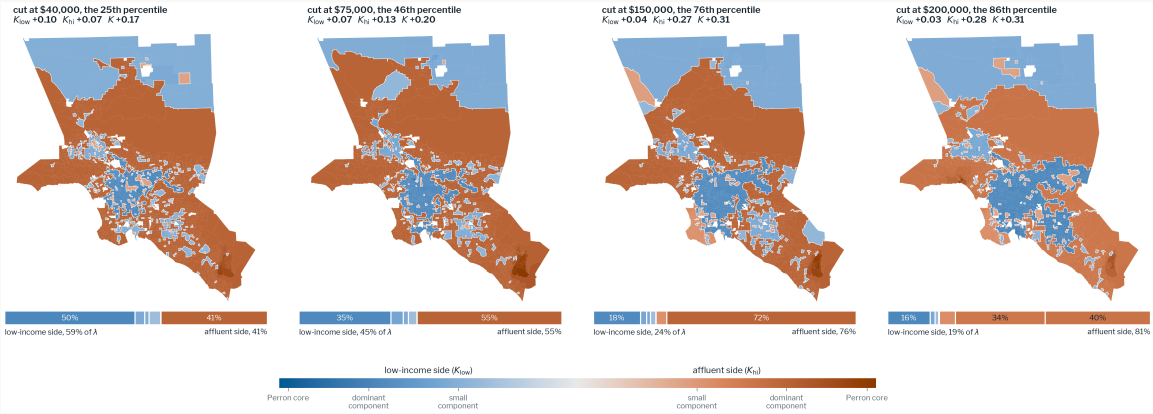


Figure 2: The concordant components behind K_{low} and K_{hi} in Los Angeles–Long Beach–Anaheim, ACS 2017–2021, at the bracket edges nearest the 25th and 50th percentiles, at the cut where $K(p)$ peaks, and at $200,000, the highest cut the brackets allow. Blue tracts over-represent households below the cut, the low-income side, and orange tracts households above it, the affluent side. Each tract's tint deepens with its component's share of its side's contribution to λ, and darkens further with its own Perron share ℓ_i of λ (Appendix E), so each component shows its core; white lines trace component boundaries. The bar beneath each map divides λ among the components: the low-income side's from the left, the affluent side's from the right, each side's largest components outermost, labeled with their share of λ where it reaches 7 percent, and its remaining components lumped beside the divide. Mainland tracts only; the Channel Islands are omitted.

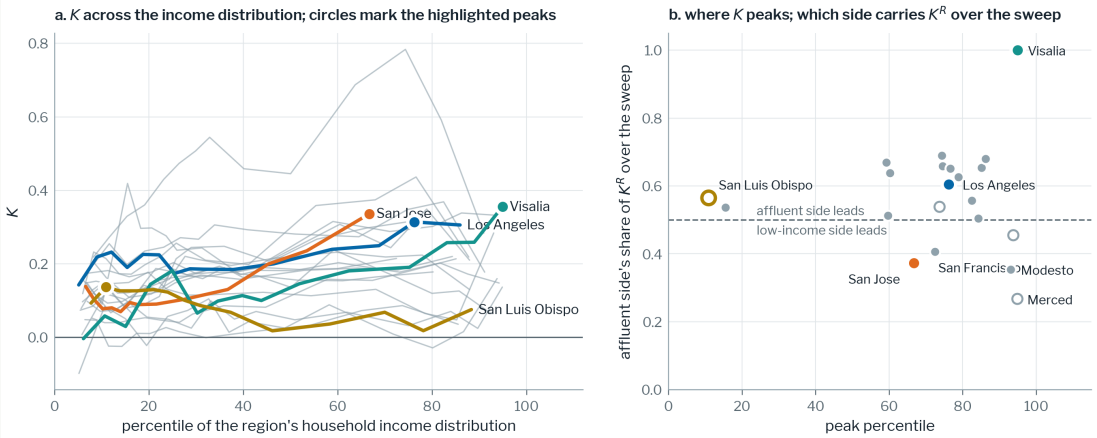


Figure 3: $K(p)$ across each region's income distribution, and where it peaks against which side carries K^R over the whole sweep. (a) $K(p)$ for the twenty-one regions, four highlighted, circles marking the highlighted curves' peaks. (b) Peak percentile against the affluent side's share of K^R over the whole sweep, K^R_{hi}/K^R; the dashed line at 0.5 divides affluent-led from low-income-led regions, and open markers are the four regions whose K^R is not significant.

6.4 Change Between 2012 and 2021

Between the 2008–2012 and 2017–2021 surveys the regional ordering on K^R held ($\rho = 0.82$), and the index fell in sixteen of the nineteen regions, by a median of 0.03, or 16 percent in relative terms (Figure 4 a). Santa Cruz, San Luis Obispo and Fresno fell most, each by about 0.11, and Stockton rose from 0.27 to 0.53, from fifth to first. Among the sixteen regions whose K^R departs from chance in both years, it fell in fourteen. By our measure, then, spatial income segregation in California's metropolitan areas generally declined over the decade.

The decline reached the whole income distribution (Figure 4 b). Over fixed bands of percentiles that both years' sweeps cover, mean K fell below the median in fifteen of the nineteen regions (median change −0.033), between the 50th and 70th percentiles in fifteen (−0.030), and above the 70th in thirteen of the eighteen whose sweeps reach it (−0.022), the smallest decline. We compare bands, and not each region's peak value, because the maximum of a noisy curve falls less than its mean when the curve flattens; the peak value itself moved from a median of 0.30 to 0.31, a stability that selection on the maximum partly produces.

The cut at which regions are most organized moved up the distribution (Figure 4 c). Its median moved only from the 74th percentile to the 75th, but the number of the nineteen regions peaking in the upper half rose from eleven to seventeen, as seven of the eight regions that peaked below the median in 2012 peaked above it in 2021; Santa Cruz moved the other way, from the 68th percentile to the 15th. Six of the seven crossed by one of two routes. In Fresno and Oxnard, concordance below the median collapsed while concordance above it held; in San Jose, Riverside, Vallejo and Stockton, concordance above the median rose. The seventh, Modesto, has a curve flat enough

that its peak moved from the 13th percentile to the 93rd, the top of its sweep, on changes of less than 0.03 in either half.

A region's peak percentile in one year does not predict its peak in the other, so the distribution of peaks across regions is more stable than any single region's. Censoring of the sweep works against this movement. The $200,000 top bracket is fixed in nominal dollars, so as incomes rose the highest available cut fell to a lower percentile of each region's distribution, in fifteen of the nineteen regions and from a median of the 93rd percentile to the 88th; six of the nineteen 2021 peaks sit at the top of their sweep, against one in 2012. Censoring of that kind pushes the 2021 peaks down, so the movement into the upper half arose despite it.

At the peak, the affluent side carried a median 0.74 of K across the nineteen regions in 2021 (0.82 among the seventeen significant in the 2021 cross-section) against 0.55 in 2012, a shift that largely follows the peaks' movement up the distribution. Over the whole sweep the split by side moved little: the affluent side led in fourteen of the nineteen regions in 2012 and sixteen in 2021, with Chico, San Luis Obispo and Santa Rosa moving to the affluent side and San Francisco to the low-income side, and the regions' shares correlate at $\rho = 0.44$ across years.

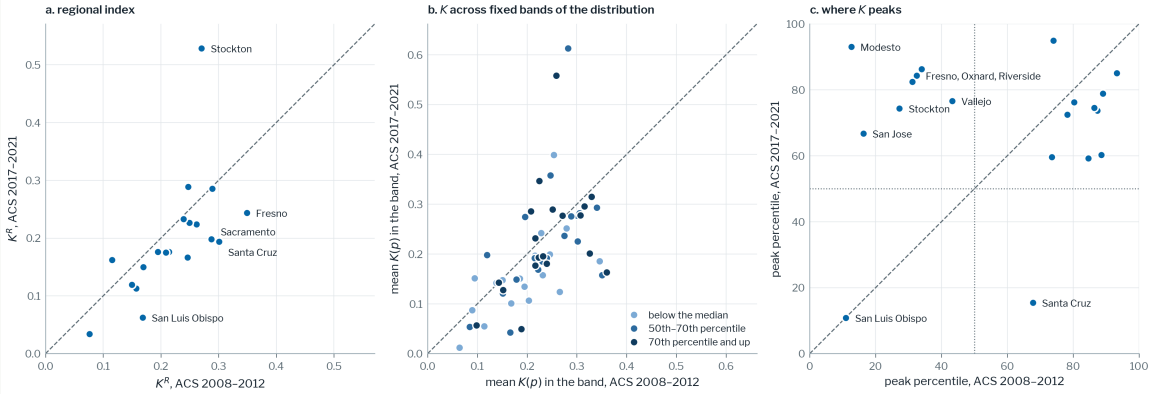


Figure 4: K^R, K across the income distribution, and where K peaks, ACS 2008–2012 against 2017–2021, for the nineteen metropolitan areas with at least fifty tracts in both. (a) K^R in each year; the dashed line is equality. (b) Mean $K(p)$ in three fixed bands of percentiles, interpolated to the cuts both years reach. (c) The peak percentile in each year; dotted lines mark the 50th percentile, so points in the upper-left quadrant are regions whose peak moved into the upper half of the distribution. Dashed lines are equality.

Put differently, concentrated affluence became the clearly dominant pattern where regions are most organized, while the spatial organization of income weakened across the distribution, least at its top. The change is a gradual and general restructuring rather than a shift at one cut, the pattern Hwang et al. (2026) describe when they argue that gentrification reorganizes segregation between race–income groups across the whole metropolitan area, beyond the neighborhoods where it occurs.

The established rank-order indices register a less consistent change. Computed on the same graphs and the same weights, Reardon and O'Sullivan's spatial information

theory index fell in ten of the nineteen regions, by a median of 10 percent, and their relative diversity index in eleven, by 1 percent, against 16 percent in sixteen for K^R; the aspatial H^R and R^R each fell in eleven, by 3 and 2 percent. Compared cut by cut, as we compare K, the two spatial indices moved by medians of less than 0.01 at their own peaks and away from them. The variance ratio $P(p)$, which measures the size of the split between tracts, already peaked in the upper half in sixteen of the nineteen regions in 2012, against eleven for K, and its median peak moved down, from the 64th percentile to the 60th.

7 Discussion and Conclusion

The segregation and polarization literatures study the same thing, the separation between portions of the income distribution, and spatial statistics supplies the frame in which their measures meet. Placed in that frame, the measures become functionals of one cross-product object, differing in how they build the attribute and in the functional they compute on the masked graph. Their disagreements become comparisons on one graph and one geography, which neither literature has occasion to make with the other's statistic.

K fills the frame's empty cell and draws on both literatures. From segregation it takes the magnitude of each unit's departure from the regional composition and, swept across thresholds, the full income distribution; from spatial polarization it takes the concept of space as connectivity, so that a region's income geography enters as a set of clusters, each measured whole by its Perron root, and splits exactly by side of the threshold. In California that combination answers, in part, the questions we began with. Spatial income segregation, as K measures it, fell between the two surveys in sixteen of nineteen regions, by a median of 16 percent, where Reardon and O'Sullivan's indices fell in only ten or eleven.

The decline reached every portion of the income distribution and was smallest at its top. The portion a region separates most sharply is an affluent minority: nearly every region is most organized at an upper cut, and there the affluent side carries most of the index. On the map, that separation takes the form of contiguous clusters, and over the whole distribution the affluent side's clustering carries more of the index in fourteen of the seventeen regions whose index departs from chance, the low-income side's only in San Francisco, San Jose and Modesto. A segregation index evaluated at one threshold, or summarized in one number, cannot separate these statements.

The index leaves several computational problems open. The first is the normalizer. The maximum concordance any rearrangement of the departures can attain is NP-hard to compute, so $\lambda^{\max}$ is defined over an explicit family of arrangements drawn from the intrinsic CAR field. The family is computed exactly and sits within a few percent of the rearrangement maximum wherever we could measure the gap (Appendix B), but it is a model-defined reference, and sharper bounds on its distance from the maximum would tighten every comparison the index supports. The second is the null. The chance

term is semi-analytic, combining percolation sweeps of each region's own graph with a closed-form bound on each cluster's strongest star and one global constant (Appendix C), and it reproduces the permutation mean closely, but the distribution of λ under random rearrangement remains unknown. A null as exact as Cohen's κ would require the polarization literature's asymptotics, which Rey (2026b) develops for component counts, to meet the randomization moments that spatial statistics derives for quadratic forms (Cliff & Ord, 1973); neither alone supplies the Perron root of a percolation cluster. Until it exists, inference runs on permutations, at the level of the region.

The third problem is the graph. K is defined on any symmetric weighted graph, but fixed-bandwidth kernels over census tracts leave dozens of disconnected pieces at every bandwidth we examined, and the normalizing family is defined only within a connected graph. Joining two islands to a mainland by their nearest edge changes little; joining dozens of pieces would impose a connectivity the kernel denies, and on a fragmented graph the components of M stop meaning clusters. A spatial-scale profile of K, the analogue of the segregation profile of Reardon et al. (2008), therefore awaits a normalizing family built for fragmented graphs. K already measures one kind of scale without a kernel, the extent of the contiguous clusters, through the spectral radius of the graph each one occupies (Section 5); a kernel profile would add the reach of the ties themselves. The fourth is local inference. The exact local split of λ places a region's index on a few tracts whether or not the map is organized, so a local statistic with power against its own null remains to be built (Appendix E).

The empirical comparison has limits of its own. Each survey year sits on its own tract boundaries, and K, like every index computed on areal units, depends on the partition. The nominal top bracket of table B19001 censors the sweep more in 2021 than in 2012. And two survey years across nineteen regions cannot separate a general trend from the idiosyncrasies of particular regions.

Two directions for future work follow directly. The first is a larger-scale study of all U.S. metropolitan areas over the full run of the American Community Survey, to establish whether the pattern in California, a general decline around a strengthening affluent peak, holds elsewhere.

The second is the extension to other networks. K needs two things: a real-valued attribute, and a connectivity graph over the units that carry it. We use income and a spatial graph, but the same construction measures concordance on any network for which the relevant attribute can be collected. On trade networks it could reveal industries or products isolated within clusters of trading partners; on peer networks, language communities closed to one another; on online social networks, sentiment exchanged within like-minded clusters; on citation networks, disciplines that cite mostly among themselves; on migration networks, wage levels that circulate within connected sets of places. On each, K answers the question a segregation index answers on a map: whether a value is confined to connected clusters, and how coherently.

On a citation network, for instance, a concept well established in one field that has not broken into a broader literature appears as a concordant cluster with few ties

leading out of it. Echenique & Fryer (2007) first measured segregation by Perron roots on friendship networks. Two requirements carry over from the spatial case. The graph must be symmetric, so directed flows such as trade, citation or migration enter as symmetrized totals until a directed form of the index exists; and it must be connected, so that its components mean clusters. On a network of individuals, moreover, each node's departure at a cut takes one of two values, so the magnitude the composition rule supplies returns only where nodes aggregate units, as tracts aggregate households; and because the semi-analytic chance term is calibrated on contiguity graphs, a social graph with heterogeneous degrees takes the permutation mean instead.

Preregistration

This study was not preregistered.

Code, Materials, and Data Availability

The tract tables used here are built from the American Community Survey's five-year estimates for 2008–2012 and 2017–2021, released publicly by the Census Bureau. All code that computes the index, draws the permutation nulls, and produces every figure and table is available for peer review at [ANONYMIZED REPOSITORY LINK: to be added before upload]. After acceptance the code and the derived tract tables move to a named public archive, and this statement will carry the link.

Institutional Review and Consent

The study analyzes publicly available aggregate census data and involves no human participants, so institutional review board approval was not required.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

References

Anselin, L. (1995). Local Indicators of Spatial Association—LISA. *Geographical Analysis*, *27*(2), 93–115. https://doi.org/10.1111/j.1538-4632.1995.tb00338.x

Besag, J. (1974). Spatial Interaction and the Statistical Analysis of Lattice Systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, *36*(2), 192–236. https://www.jstor.org/stable/2984812

Besag, J., & Kooperberg, C. (1995). On Conditional and Intrinsic Autoregression. *Biometrika*, *82*(4), 733–746. https://doi.org/10.2307/2337341

Bischoff, K., & Reardon, S. F. (2014). Residential segregation by income, 1970–2009. In J. R. Logan (Ed.), *Diversity and disparities: America enters a new century* (pp. 208–233). Russell Sage Foundation.

Carrington, W. J., & Troske, K. R. (1997). On measuring segregation in samples with small units. *Journal of Business & Economic Statistics*, *15*(4), 402–409. https://doi.org/10.1080/07350015.1997.10524720

Cliff, A. D., & Ord, J. K. (1973). *Spatial autocorrelation*. Pion.

Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, *20*(1), 37–46. https://doi.org/10.1177/001316446002000104

Dawkins, C. J. (2007). Space and the measurement of income segregation. *Journal of Regional Science*, *47*(2), 255–272.

de Jong, P., Sprenger, C., & van Veen, F. (1984). On Extreme Values of Moran's I and Geary's c. *Geographical Analysis*, *16*(1), 17–24. https://doi.org/10.1111/j.1538-4632.1984.tb00797.x

Duclos, J.-Y., Esteban, J., & Ray, D. (2004). Polarization: Concepts, measurement, estimation. *Econometrica*, *72*(6), 1737–1772. https://doi.org/10.1111/j.1468-0262.2004.00552.x

Echenique, F., & Fryer, R. G. (2007). A measure of segregation based on social interactions. *The Quarterly Journal of Economics*, *122*(2), 441–485. https://doi.org/10.1162/qjec.122.2.441

Esteban, J.-M., & Ray, D. (1994). On the measurement of polarization. *Econometrica*, *62*(4), 819–851. https://doi.org/10.2307/2951734

Getis, A. (1991). Spatial interaction and spatial autocorrelation: A cross-product approach. *Environment and Planning A: Economy and Space*, *23*(9), 1269–1277. https://doi.org/10.1068/a231269

Getis, A., & Ord, J. K. (1992). The analysis of spatial association by use of distance statistics. *Geographical Analysis*, *24*(3), 189–206. https://doi.org/10.1111/j.1538-4632.1992.tb00261.x

Hubert, L. J., Golledge, R. G., & Costanzo, C. M. (1981). Generalized procedures for evaluating spatial autocorrelation. *Geographical Analysis*, *13*(3), 224–233. https://doi.org/10.1111/j.1538-4632.1981.tb00731.x

Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, *2*(1), 193–218. https://doi.org/10.1007/BF01908075

Hwang, J., Roberto, E., Stephens, F. I., & Rugh, J. S. (2026). Integrating or segregating? Gentrification, the changing city-suburb divide, and black exceptionalism. *Social Forces*, soag102. https://doi.org/10.1093/sf/soag102

James, D. R., & Taeuber, K. E. (1985). Measures of segregation. *Sociological Methodology*, *15*, 1–32. https://doi.org/10.2307/270845

Jargowsky, P. A. (1996). Take the money and run: Economic segregation in U.S. Metropolitan areas. *American Sociological Review*, *61*(6), 984–998. https://doi.org/10.2307/2096304

Logan, J. R., Foster, A., Ke, J., & Li, F. (2018). The uptick in income segregation: Real trend or random sampling variation? *American Journal of Sociology*, *124*(1), 185–222. https://doi.org/10.1086/697528

Mantel, N. (1967). The detection of disease clustering and a generalized regression approach. *Cancer Research*, *27*(2), 209–220. https://www.ncbi.nlm.nih.gov/pubmed/6018555

Massey, D. S., & Denton, N. A. (1988). The dimensions of residential segregation. *Social Forces*, *67*(2), 281–315. https://doi.org/10.2307/2579183

Ord, J. K., & Getis, A. (1995). Local spatial autocorrelation statistics: Distributional issues and an application. *Geographical Analysis*.

Owens, A., Reardon, S. F., & Jencks, C. (2016). Income segregation between schools and school districts. *American Educational Research Journal*, *53*(4), 1159–1197. https://doi.org/10.3102/0002831216652722

Reardon, S. F. (2011). *Measures of Income Segregation*. https://inequality.stanford.edu/sites/default/files/media/_media/working_papers/reardon_measures-income-seg.pdf

Reardon, S. F., & Bischoff, K. (2011). Income Inequality and Income Segregation. *American Journal of Sociology*, *116*(4), 1092–1153. https://doi.org/10.1086/657114

Reardon, S. F., Bischoff, K., Owens, A., & Townsend, J. B. (2018). Has Income Segregation Really Increased? Bias and Bias Correction in Sample-Based Segregation Estimates. *Demography*, *55*(6), 2129–2160. https://doi.org/10.1007/s13524-018-0721-4

Reardon, S. F., & Firebaugh, G. (2002). Measures of multigroup segregation. *Sociological Methodology*, *32*(1), 33–67. https://doi.org/10.1111/1467-9531.00110

Reardon, S. F., Matthews, S. A., O'Sullivan, D., Lee, B. A., Firebaugh, G., Farrell, C. R., & Bischoff, K. (2008). The geographic scale of Metropolitan racial segregation. *Demography*, *45*(3), 489–514. https://doi.org/10.1353/dem.0.0019

Reardon, S. F., & O'Sullivan, D. (2004). Measures of Spatial Segregation. *Sociological Methodology*, *34*(1), 121–162. https://doi.org/10.1111/j.0081-1750.2004.00150.x

Rey, S. J. (2026a). Mind the gap and the map: Measuring distributional and spatial income polarisation. *Spatial Economic Analysis*, *21*(1), 41–60. https://doi.org/10.1080/17421772.2025.2606783

Rey, S. J. (2026b). On the distribution of spatial polarization indices. *Geographical Analysis*, *59*(1). https://doi.org/10.1111/gean.70064

Rey, S. J., & Folch, D. C. (2011). Impact of spatial effects on income segregation indices. *Computers, Environment and Urban Systems*, *35*(6), 431–441. https://doi.org/10.1016/j.compenvurbsys.2011.07.008

Sampson, R. J., Mare, R. D., & Perkins, K. L. (2015). Achieving the Middle Ground in an Age of Concentrated Extremes. *The ANNALS of the American Academy of Political and Social Science*, *660*(1), 156–174. https://doi.org/10.1177/0002716215576117

Trounstine, J. (2020). The geography of inequality: How land use regulation produces segregation. *American Political Science Review*, *114*(2), 443–455. https://doi.org/10.1017/S0003055419000844

White, M. J. (1986). Segregation and diversity measures in population distribution. *Population Index*, *52*(2), 198–221. https://doi.org/10.2307/3644339

Wolfson, M. C. (1997). Divergent inequalities: Theory and empirical results. *Review of Income and Wealth*, *43*(4), 401–421. https://doi.org/10.1111/j.1475-4991.1997.tb00233.x

Zhang, X., & Kanbur, R. (2001). What difference do polarisation measures make? An application to China. *Journal of Development Studies*, *37*, 85–98.

A Synthetic Geographies

The synthetic geographies of Figure 1 cover a disc of radius 25 kilometers with 360 hexagonal units at H3 resolution 7, on queen contiguity, about the size of San Jose's 420 tracts. Households belong to three latent income classes, lognormal around medians of $25,000, $70,000 and $220,000, with regional shares near 0.25, 0.55 and 0.20. Each unit holds 1,500 households drawn from its own mixture of classes and binned to the sixteen brackets of table B19001, so the sweep uses the same cuts as the California data, with sampling noise. The scenarios move only the geography of the class mixture: the affluent class concentrated in three compact enclaves; the low-income class concentrated in three pockets; both, apart; and the enclave scenario's compositions with the enclave units relocated to isolated single units, so the multiset of unit compositions, and with it $P(p)$, is unchanged. The chance term is the 199-draw permutation mean throughout.

The shape comparison in Section 5 holds one multiset of unit compositions fixed and arranges the affluent units in five shapes: a compact blob, a wedge, a ring of hexagons, a band, and scattered single units, each at 24 units, with the band one unit wide, and at 54, with the band about two units wide, so that P is identical within each size. The ratios reported are each measure's value on a shape divided by its value on the blob, with $\tilde{R}$ computed on environments of one and two rings of neighbors.

B The Normalizing Family

B.1 The Rearrangement Maximum Is NP-Hard

On a d-regular graph the standardized weights are $W = A/d$. Take the multiset of departures $\{+1^{(m)}, -\varepsilon^{(n-m)}\}$, up to a common scale, which leaves the maximizing arrangement unchanged. The negative class contributes $O(\varepsilon^2)$ to λ, and the positive class at most

$m(m-1)/(dn)$, with equality exactly when the positive units form an m-clique, so the maximum over rearrangements decides whether the graph contains an m-clique. The multiset is realizable as composition departures, with shares in $[0, 1]$ and population-weighted mean zero, so neither the centering nor the household weighting blocks the reduction. Enumerating every arrangement on forty random 4-regular graphs of ten units with $m = 4$ puts the maximum at exactly $m(m-1)/(dn) = 0.300$ on every graph holding a 4-clique and at most 0.854 of that value on the rest. The index is defined on any symmetric graph, so hardness over regular graphs of unbounded degree rules out a formula; whether the maximum remains hard on contiguity graphs, whose cliques are bounded by the number of units meeting at a point, is open. A heuristic search is no substitute: its value is a lower bound that any better search would raise, so an index normalized by it moves whenever the search does.

B.2 The Family

For a connected graph with adjacency A and Laplacian $L = D - A$, the intrinsic CAR field conditioned on a high value at unit c and a low value at unit c' has conditional mean g, the harmonic potential with $Lg = 0$ away from the two units, $g_c = 1$ and $g_{c'} = -1$: the potential across the graph when current is driven from c to c'. Each member of the family sorts the population-weighted departures $z^{(t)}$ along g, placing the largest positive departures where g is highest and the largest negative departures where it is lowest. The family ranges c over every unit, c' over the twenty units at the largest effective resistance from c, and both labellings, z and $-z$; values of g within 10^{-10} of one another are ties, broken by unit index, so that rounding in the solve cannot decide the sort. The family has no scale parameter. $\lambda^{\max}$ is the largest λ over its members, a deterministic function of the graph and the sorted departures, and one denominator serves the observed map and every rearrangement of it.

B.3 Exact Computation

A side-level Collatz–Wielandt bound, $\lambda \le \sum_s (n_s/n) \max_{i \in s} (Mx)_i / x_i$ over the two sides s for any positive x, with the side sizes n_s taken from each arrangement, bounds every member's λ from a few sparse matrix products. All but a few dozen members fall below the best exactly evaluated member and are excluded; the remaining members are evaluated exactly, best bound first, until no remaining bound can win. The potentials come from sparse Cholesky solves of the grounded Laplacian, with no eigendecomposition, so the value is reproducible to the last digit across solvers and thread counts. Against exhaustive evaluation of every member, the pruned maximum agrees to 10^{-9}. A Los Angeles cell, the largest, takes about five seconds.

B.4 Distance from the Rearrangement Maximum

Where the rearrangement maximum can be computed, by enumerating every permutation of the departures on 118 connected neighborhoods of eight to eleven California tracts, the family lies within 0.5 percent of it in 91, within 1 percent in 103, and 4.0

percent short at worst. On twelve metropolitan areas of 50 to 173 tracts, at cuts near the 25th and 75th percentiles, simulated annealing from the family's best arrangement and from random ones exceeds the family in 21 of 24 cells, by a median of 0.4 percent and at most 4.1 percent; normalized by the best arrangement annealing finds, K moves by at most 4.8 percent of its value, and by less than 1 percent in 17 of the 24 cells. The shortfall is unrelated to the cut, the size of the minority side, or the size of the graph. No observed map exceeds the family.

C The Chance Term

Under random rearrangement of the units' (composition, household count) pairs, the units on each side of a threshold form a uniformly random subset of fixed size, so the component structure of M is site percolation on the region's graph, a property of the graph alone. We estimate the component-size law of each graph from 32 seeded Newman–Ziff percolation sweeps, computed once per graph and shared across thresholds.

Each component then contributes through its strongest same-side star. Each block M_k of M is symmetric and non-negative, so $e_i' M_k^2 e_i \le \lambda_{\max}(M_k^2) = \lambda_k^2$, and the largest diagonal entry of M_k^2 bounds the component's Perron root from below, $\lambda_k \ge \left(\max_{i \in C_k} \sum_j M_{ij}^2\right)^{1/2}$. The star sum $\sum_j M_{ij}^2$ has moments under the null that follow in closed form from W, the share of units on that side, and the moments of the squared departures. Adjacent stars share terms, so a component of s units carries about $s/2$ independent stars, and the expected maximum over them follows from extreme-value theory. One global constant, $C_\star = 1.1934$, scales the expected strongest star to the expected Perron root. It is the median ratio of the two over the same 300 California cells on which we evaluate the term below, so the synthetic geographies are its out-of-sample check. The chance term is therefore semi-analytic: the percolation law comes from sweeps of the graph, the star bound is closed form, and no term permutes the attribute.

Computed with it, K differs from K computed with a 199-permutation estimate of $\mathbb{E}[\lambda]$ by a median of 0.005 across the 300 threshold-region cells, by less than 0.02 in 97 percent of them, and places every region's peak at the same threshold. On the synthetic hexagonal geographies of Appendix A it matches the permutation mean to within 0.014 in K over 91 cells. The term is calibrated on irregular contiguity graphs with heterogeneous departures; on regular lattices with equal departures, the fixtures of Appendix D, we use the permutation mean.

D Desirable Properties

Table 3: K against the criteria of Reardon & O'Sullivan (2004), measured on a 12×12 lattice with the chance-corrected index of Equation 12, the permutation mean as its chance term.

	Criterion	K
1	Scale interpretability	holds in chance-anchored form, through the correction and the ceiling
2	Arbitrary boundary independence	fails
3	Location equivalence	holds to a vanishing residual
4	Population density invariance	holds
5	Composition invariance	fails
6	Transfers and exchanges	fails
7	Additive spatial decomposability	fails; replaced by the exact split by side
8	Additive grouping decomposability	does not apply

Scale interpretability holds in modified form, anchored at chance, and only because of the correction and the ceiling. M is non-negative, so the uncorrected ratio $\lambda/\lambda^{\max}$ cannot fall below zero: a checkerboard scores zero while a random map scores its chance floor, above it. Corrected, two perfectly separated blocks score 1, random maps score 0 on average, and the rook checkerboard scores −1.43. Density invariance holds exactly: scaling every unit's household count by a constant leaves the shares unchanged and scales λ, $\lambda^{\max}$ and $\mathbb{E}[\lambda]$ alike, since each is homogeneous of degree one in the household counts. Location equivalence holds to a vanishing residual: merging two equivalent units moves K by 0.0037, 0.15 percent of the range the lattice attains.

Boundary independence fails, as it does for every measure computed on aggregated units: merging the lattice's cells two by two and three by three moves K from 0.106 to 0.370 and 0.553. Composition invariance fails as it does for the aspatial parent: scaling one group's count by a constant at every unit, as James & Taeuber (1985) define the criterion, moves K from 0.122 to 0.092 as the regional share runs from 0.24 to 0.89, which places K with the exposure measures. The exchange criterion, which Reardon & O'Sullivan (2004) state on local environments, fails in a minority of exchanges, as it does for every spatial measure they describe in the general spatial case. A Perron root answers to a whole component, so an exchange that moves two units toward the regional composition can also join two concordant clusters that had been separate, and the second effect occasionally outweighs the first. Additive spatial decomposability fails for any measure built on an eigenvalue, since a Perron root is not additive over a partition of the units; the exact split by side, and the exact division of λ among units in

Appendix E, take its place. K is a two-group index at each threshold, and the sweep stands in for multigroup treatment, so grouping decomposability does not apply.

E The Local Decomposition

The Perron vector v_k of each component, normalized to unit length, divides that component's root among its units, $\lambda_k = \sum_{i \in C_k} \lambda_k v_{ki}^2$. Writing $\ell_i = |C_k| \lambda_k v_{ki}^2$ for unit i in component C_k, the mean of ℓ_i over units is exactly λ, and the local index $K_i = (\ell_i - \mathbb{E}[\lambda])/(\max(\lambda^{\max}, \lambda) - \mathbb{E}[\lambda])$ has mean K, the property that defines a local indicator of spatial association (Anselin, 1995). The local index concentrates on a few tracts whether or not the map is organized, because a Perron vector concentrates on the core of its component in random arrangements as in observed ones. Tested tract by tract with permutations and a false discovery rate of 0.10 within each region, 15 of California's 8,620 metropolitan tracts are significant, and the concentration of the local index within a region never exceeds its own null. We therefore read the local index descriptively and conduct inference at the level of the region.