Session Number: 7B
Session Title: Measurement of Segregation: New Directions and Results
Session Organizers: Yves Flückiger and Jacques Silber
Session Chair: Jacques Silber

Paper Prepared for the 29th General Conference of
The International Association for Research in Income and Wealth

Joensuu, Finland, August 20 – 26, 2006

A New Approach to Measuring Socio-Spatial Economic Segregation

by

sean f. reardon
Stanford University

Glenn Firebaugh
Pennsylvania State University

David O'Sullivan
University of Auckland

Stephen Matthews
Pennsylvania State University

Revised, August, 2006

For additional information please contact:

sean f. reardon
520 Galvez Mall, #526,
Stanford University,
Stanford, CA 94305
sean.reardon@stanford.edu
+1 650.723.9931 (fax)
+1 650.736.8517 (voice)

This paper is posted on the following websites: http://www.iariw.org
http://www.pop.psu.edu/mss

Abstract

In this paper we propose two measures of income segregation, the *rank-order information theory index* (H_R), and the *rank-order variation ratio index* (H_R). We propose spatially-explicit versions of these indices as well. These indices have several appealing features that remedy flaws in existing measures. First, they are relatively easy to compute, since they require simply computing a series of pairwise segregation values using existing measures of segregation, fitting a polynomial regression line using WLS, and then computing a linear combination of the estimated parameters. Second, the measures are easily adapted to account for spatial proximity, following the approach of Reardon and O'Sullivan (2004). Third, the measures are largely insensitive to the set of income thresholds that define the categories in which income is reported, so long as the thresholds span most of the range of income percentiles. As a result, they do not require us to make assumptions about the shape of income distributions. Fourth, the measures are insensitive to rank-preserving changes in income, since the measures are based on the ranks of incomes rather than their numerical values. Finally, the indices can be interpreted in a variety of equivalent ways that illustrate their correspondence with standard notions of segregation.

Introduction

In any city in the world, even a cursory inspection—by any non-sociologist, non-economist, or non-geographer—of residential patterns would indicate the presences of some degree of residential segregation by income and wealth—there are some neighborhoods populated primarily by families with above-average income and wealth, and other neighborhoods populated primarily by families with below average income and wealth. For the sociologist or economist, however, it is not enough to merely note the presence of such residential sorting; we desire as well to quantify it. Virtually any interesting question regarding the causes, patterns, and consequences of such residential segregation requires that we measure it—and measure it in a way that makes comparisons across places and times meaningful.

Surprisingly, the set of tools available to scholars for measuring spatial economic segregation is relatively limited. Our goal in this paper is to develop an approach to measuring spatial economic segregation that is intuitively meaningful, easy to compute, and allows for comparisons across place and time. To achieve this, we develop several measures of segregation along an ordinal dimension (since income data and other socioeconomic data such as educational attainment and occupational status are often reported using ordered categories). Although we initially develop 'aspatial' versions of these measures, we show how they can be easily adapted to take into account the spatial or social proximity of individuals by using the approach outlined by Reardon and O'Sullivan (2004). Our initial ordinal segregation measures, however, are sensitive to changes in the overall distribution of the quantity to be measured as well as to changes in how the ordered categories are defined, both of which make comparisons difficult. To remedy this, we propose a method of adjusting the measures so that they are invariant to such changes and allow meaningful comparisons across places. We then illustrate our method using data on residential segregation by household income in a set of metropolitan areas in the United States.

The paper proceeds as follows. Section 1 reviews existing methods of measuring economic segregation. Section 2 develops two simple measures of ordinal segregation that can be used to measure economic segregation, and notes that both are flawed. Section 3 describes an approach that remedies flaws in these two measures. In section 4, we apply our measures to describe residential segregation by household income in a set of metropolitan areas in the United States in 2000. Section 5 concludes.

1. Existing Measures of Economic Segregation

Prior research on economic segregation has relied on several general approaches for characterizing the extent to which individuals of different socioeconomic characteristics are unevenly distributed throughout a region. Most of this research has been concerned with income segregation, rather than segregation by wealth, largely because income data are far more readily available. This research comes primarily from three different disciplinary perspectives—sociology, economics, and geography—each of which faces the same set of measurement issues.

In general, income data are reported categorically, as counts within each organizational unit (e.g., census tract) of households, families, or individuals falling in each of a set of mutually exclusive and exhaustive ordered income categories. Each of these income categories is defined by a pair of upper and lower income bounds (except for the two extreme categories, which are each unbounded on one side)[1]. For example, in the 2000 U.S. census, household income is categorized by 16 annual income categories, ranging from "less than $10,000," "$10,000-$14,999", $15,000-$19,000, etc., through "$150,000-$199,999", and finally "$200,000 or more." As a result, the measurement of income segregation is hampered by the fact that we generally do not know individuals' exact incomes (and so we lack full information on the income distribution overall or in any one

[1] Since negative income is possible, the lowest income category is not bounded by 0.

organizational unit). Moreover, although the income thresholds that define the ordered income categories are designed to span most of the range of incomes, they are nonetheless relatively arbitrary and change over time, so that any segregation measure that relies on them should be insensitive to the definition of the thresholds if the measure is to be useful for comparative purposes.

Category-Based Measures of Economic Segregation

By far the most common method of measuring income segregation used in existing research has been to divide the population into two categories, based on some chosen income threshold (a wide range of thresholds are used in extant research). Segregation between these two groups (those above and those below the chosen threshold) is computed using any conventional two-group segregation measure, such as the dissimilarity index. Examples of this approach are found in the literature in sociology (Fong, 2000; Massey, 1996; Massey & Eggers, 1993; Massey & Fisher, 2003), urban planning (Coulton, Chow, Wang, & Su, 1996; Pendall & Carruthers, 2003), and economics (Jenkins, Micklewright, & Schnepf, 2006; Waitzman & Smith, 1998).

Although the primary advantage of this approach is its simplicity, its shortcomings are several and obvious. First, dichotomizing the income distribution discards an enormous amount of information. Even if we don't know individuals' exact incomes, the 16 income categories reported in the U.S. census, for example, contain far more information than any dichotomized version. Second, the results of such an approach may depend on the choice of a threshold—segregation between the very poor and everyone else may not be the same (and generally is not the same) as segregation between the very rich and everyone else.

Several variants of this approach have been used. Massey and Fischer (2003), for example, compute segregation between poor and affluent households (ignoring the middle-class) to better capture the separation of the extremes of the income distribution. A second variant of the categorical approach to measuring income segregation is to compute the two-group segregation for many or all possible pairs of income categories, and then to construct some average or summary measure of these multiple pairwise indices (Farley, 1977; Massey & Eggers, 1990; Telles, 1995). Third, rather than dichotomize the income distribution, Fong and Shibuya (2000) and Telles (1995) compute segregation among multiple income category groups using the Theil information theory index of segregation (Theil, 1972). This approach, however, uses an index designed to measure segregation among a set of unordered groups (such as racial groups) to measure segregation among a set of ordered groups (income groups), and so is insensitive to the inherently ordinal nature of income segregation. Finally, Meng, Hall, and Roberts (2006) measure the segregation among multiple ordered income groups using an approach that explicitly accounts for the ordered nature of the categories by weighting the segregation between different groups by some measure of the 'social distance' between the groups. While many of these variants have some advantage over simply dichotomizing the income distribution, each shares with all categorical income segregation approaches the fundamental flaw that they are sensitive to the number and location of the thresholds used to define income categories, confounding the possibility of making meaningful comparisons across places and times.

Variation-Ratio Measures of Economic Segregation

A second approach to measuring income segregation defines segregation as a ratio of the between-neighborhood variation in mean income or wealth to the total population variation in income or wealth. Some measures derived from this approach use the variance of incomes as the measure of income variation (Davidoff, 2005; Wheeler, 2006; Wheeler & La Jeunesse, 2006). Such measures have an interpretation analogous to the R^2 statistic from a regression of individual incomes on a set of neighborhood dummy variables. Other such measures use a measure of variation other than the variance as the chosen measure of variation. Jargowsky (1996; , 1997) for example, defines income segregation as the ratio of the between-unit (e.g., between-tract) standard deviation of income to the overall regional income standard deviation. Ioannides (2004) uses the ratio of the variance of log incomes; Hardman & Ioannides (2004) use the ratio of within-neighborhood to overall coefficients of variation of income; Ioannides and Seslen (2002) use the ratio of Bourguignon's population-weighted decomposable inequality index to measure both income and wealth segregation. The relative merits and flaws of the choice of inequality or variation measure used to construct the various ratio-based indices have not been fully investigated.

In principle, measures based on this approach use full information on the income distribution at each location, but since exact income distribution data are generally not available, they generally must rely in part on the estimation of parameters describing the overall income distribution (see, e.g., Jargowsky, 1996; Wheeler & La Jeunesse, 2006). This estimation, in turn, may be very sensitive to assumptions about the income levels of individuals in the top income category. To the extent that the required parameters (e.g., variance) of the income distribution can be estimated well from the reported counts by income category, however, variation ratio approaches have considerably more appeal than existing approaches rely on computing pairwise segregation between groups defined by one or more income thresholds. They use (in theory) complete information on the income distribution; they do not rely on arbitrary threshold choices; and at least some such measures are invariant to certain types of changes in the income distribution (e.g., Jargowsky's *NSI* measure is invariant under shape-preserving changes in the income distribution).

One interesting variant of the variation ratio approach is developed by Watson (2006), who measures income segregation using the Centile Gap Index (*CGI*), which is defined as one minus the ratio of a measure of the within-neighborhood variation in income *percentile ranks* to the overall variation in income percentile ranks. Specifically, the *CGI* measures within-neighborhood variation in percentile ranks as the mean absolute deviation of households' income percentile from the percentile rank of their neighborhood median. The *CGI*, because it is based on variation in percentile ranks rather than income levels, is insensitive to any rank-preserving changes in the income distribution, a feature that all ratio measures lack.[2]

<u>Spatial Autocorrelation Measures of Economic Segregation</u>

In general, most proposed measures of income segregation are *aspatial*—that is, do not account for the spatial proximity of individuals/households, except insofar as spatial proximity is accounted for by census or administrative area boundaries. Jargowsky and Kim (2005), however, describe a spatial version of Jargowsky's income segregation measure, but theirs is an exception.

A third approach to measuring income segregation derives from the geographical notion of spatial autocorrelation. In this approach, which explicitly accounts for the spatial patterning of households, segregation is conceived as the extent to which households near one another have more similar incomes than those that are farther from one another. Although several such measures have been suggested (Chakravorty, 1996; Dawkins, forthcoming), this approach to measuring income segregation is the least well-developed.

In sum, while a wide range of measures have been used to describe income segregation, several key flaws plague existing measures. Measures based on computing categorical segregation indices among income categories, while widely used because of their ease of computation, are

[2] While Watson's *CGI* has some appeal because of its insensitivity to rank-preserving changes in income, it has a subtle flaw. It is insensitive to redistributions of individuals among neighborhoods that do not affect the median income in each neighborhood. As a result, if we have a region consisting of two neighborhoods with identical income distributions (so that *CGI*=0), and we rearrange households so that one neighborhood consists of the households in the first and third quartiles of each prior neighborhood and the other consists of the households in the second and fourth quartiles of each prior neighborhood, the CGI will be unchanged (*CGI*=0), despite the fact that we have created an uneven distribution of households among neighborhoods, such that the two neighborhoods now have different mean incomes.

inherently sensitive to the definition of income thresholds used to define income categories. As a result, changes in either the choice of thresholds (as occurs between censuses in the U.S.) or differences in income distributions (either regional or temporal) may affect measured segregation, even in the absence of any change in the location and relative income levels of households. If a measure of income segregation is to be based on discrete income categories, we would like it to be insensitive to the choice of thresholds, a feature which no existing categorical income segregation measure possesses.

Measures based on ratios of income variation within and among locations, in contrast, do not depend on the definition of income categories, at least in principle. In practice, however, the distributional parameters used in such measures must be estimated from the categorical income data generally reported, and so may be sensitive not only to the definition of categories, but also to assumptions about the shape of the income distribution, particularly for the highest-earning category, which has no upper bound. Moreover, it is not clear how sensitive the measurement of income segregation is to the choice of a measure of income spread (variance, standard deviation, variance of log income, coefficient of variation, Bourguignon inequality, etc.); nor is it clear on what basis one should choose among these. The standard deviation or variance ratio measures (as used, for example, by Jargowsky, 1996, 1997; Wheeler, 2006; Wheeler & La Jeunesse, 2006) are insensitive to shape-preserving changes in the income distribution (changes that shift the mean and/or multiply all incomes by a constant), but measures based on other parameters, such as the variance of logged incomes are not (e.g., the variance of logged income is insensitive to constant multiplicative changes in incomes, but is sensitive to changes that add a constant to all incomes).

Measures based on spatial autocorrelation are the least well-developed set of measures of income segregation. Although they have the advantage of being explicitly spatial, they generally rely on areal unit boundary definitions, and so are subject to the modifiable areal unit problem (MAUP)

(Openshaw, 1984). Moreover, even in the absence of MAUP issues, the measurement properties of spatial autocorrelation measures of income segregation are not well understood.

An ideal income segregation measure would have several key features. First, it would be threshold invariant. Second, it would be invariant under specified distributional changes, although the specific type of distributional invariance desired may depend on the application. At a minimum, we might wish the measure to be scale invariant with respect to the income distribution. Such a property ensures, for example, that doubling each household's income does not affect measured segregation. We might also wish the measure to be invariant under a change in the mean income—as would occur if each household income increased by a constant amount. A stronger invariance property is invariance under rank-preserving changes in income. This property implies the other two, but also requires that segregation not change when incomes change in a way that does not affect the rank-ordering of households in the income distribution. Watson (2006), for example, notes the advantage of measures with this property when examining income sorting processes. A third key feature of a segregation measure is that it account for the spatial patterning of households, preferably in a manner that renders it insensitive to MAUP issues. In the remainder of this paper, we develop measures of income segregation that meet these three criteria—the measures are threshold-invariant, invariant under rank-preserving changes in income, and MAUP-free.

2. Measuring Segregation by an Ordinal Category

Because income data are generally reported as counts by ordered income category, we begin by describing an approach to measuring segregation among groups defined by ordinal categories. Reardon and Firebaugh (2002) and Watson (2006) suggest that one way of constructing a segregation measure is to think of it as a form of variance decomposition, where segregation is the proportion of the total variation in a population that is due to differences in population composition of different organizational units (e.g., schools or census tracts). This approach is akin to the variation-ratio segregation measures described above. For an unordered categorical variable (such as race), population variation is measured by diversity or entropy, but for an ordered (ordinal or continuous) variable, variation is typically measured using some index of the spread of the distribution (e.g., variance, in the case of an interval-scaled variable). Following Reardon and Firebaugh (2002, Eq. 9, p. 45), if we have a suitable measure of variation v, we can define $S(v)$—a segregation measure based on the variation measure v—as follows:

$$S(v) = 1 - \frac{\bar{v}_j}{v} = \sum_{j=1}^{J} \frac{t_j}{Tv}\left(v - v_j\right), \tag{1}$$

where j indexes organizational units t_j and v_j are the population count and variation in unit j, $\bar{v}_j$ is the weighted average within-unit variation over the region, and T is the total population of the region. Note that if v is the variance of a continuous variable x, then $S(v)$ is equivalent to η^2 (or the R^2 from a regression of x on a set dummy variables for organizational units).

One way to construct a measure of ordinal segregation, then, is to define a suitable measure of ordinal variation (we define what we mean by 'suitable' below), and then use it to construct a segregation measure as above.

Ordinal variation

Measuring the variation in a population of a quantity measured with an ordinal variable requires us to define what we mean by variation. For an ordinal variable x that can take on any of K ordered categories 1, 2, …, K, we define variation as having a maximum (which we can normalize to equal 1) when half the population has $x=1$ and half has $x=K$. Variation is at a minimum (defined as

0) when all observations have $x=k$ for some $k \in 1, 2, \ldots, K$. Measuring ordinal variation then amounts to measuring how close the distribution of x is to these minimum and maximum variation states.

It is useful to express the distribution of values of x in a sample as a [K-1]-tuple of cumulative proportions, $C=(c_1, c_2, \ldots, c_{K-1})$, where c_k is the cumulative proportion of the sample with values of X in category k or below (note that $c_K=1$ by definition, so is not needed to characterize the distribution of x). Note that a distribution of x has maximum variation at $C_0=(½, ½, ½, \ldots, ½)$, corresponding to the case where half the population has the lowest possible value and half has the highest possible value of x. Moreover, note that there are K possible distributions of x such that there is no variation in x, corresponding to the K cases of the pattern $C=(0,0,\ldots 0,0,1,1,\ldots 1,1)$ (where all observations have the same value of $x=k$, so that $c_j=0$ for $j<k$ and $c_j=1$ for all $j \geq k$).

Blair and Lacy (1996) suggest that it is helpful to think of C as a point in [K-1]-space, which leads to the insight that variation can be measured as an inverse function of the distance from C to C_0, the point of maximum variation (it is easier to define variation in terms of the distance from the single point of maximum variation rather than from one of the K points of zero variation). Alternately, we can think of C as describing a cumulative density function of an ordinal variable, where variation is measured as an inverse function of the average distance of C from the line $y=½$. This suggests a general form of a variation measure as:

$$v(x) = \frac{1}{K-1}\sum_{i=1}^{K-1} f(c_i), \qquad (2)$$

where $f(c)$ is maximized at $f(½)=1$ and minimized on the interval [0,1] at $f(0)=f(1)=0$. Three such possible functions f are given by (where we define $0\log_2(1/0)=0$):

$$
\begin{aligned}
f(c) &= -\left[c\log_2 c + (1-c)\log_2(1-c)\right] \\
f(c) &= 4c(1-c) \\
f(c) &= 1-|2c-1|
\end{aligned}
\quad . \tag{3}
$$

Here we consider two[3] specific measures of ordinal variation, E_O and I_O, defined as follows, where c_k is the cumulative proportion of observations in categories 1 through k:

$$E_O = \frac{-1}{K-1}\sum_{k=1}^{K-1}\left[c_k \log_2(c_k) + (1-c_k)\log_2(1-c_k)\right] \tag{4}$$

$$I_O = \frac{1}{K-1}\sum_{k=1}^{K-1} 4c_k(1-c_k) \tag{5}$$

The first of these, E_O, is an index we call the *index of ordinal entropy*; the second, I_O, is the *index of ordinal variation* (Kvålseth, 1995b).[4] These indices can be seen as measures of the average deviation of c_1 c_2,… c_{K-1} from their values in a state of no variation (when the c_k each equal zero or one), where 'deviation' is measured by a metric defined by f.

Both E_O and I_O equal their maximum value, 1, if and only if the distribution of x is such that $x=1$ for half the observations and $x=K$ for the other half (corresponding to the cumulative proportion vector $C=(½, ½, ½,…, ½, 1)$), and both equal their minimum value of 0 if and only if all observations have the same value of x (corresponding to a cumulative proportion vector $C=(0,0,…0,0,1,1,…,1,1)$). Note that, in the special case where K=2, E_O and I_O are identical to the

[3] We do not consider a measure of ordinal variation based on the third function, $f(c)=1-|2c-1|$, because it does not yield a satisfactory index of ordinal segregation. In fact, substituting $f(c)=1-|2c-1|$ into Equation (2) and the substituting the resulting equation into (1) yields Watson's *CGI*, which, as we note above, is flawed.

[4] Several alternate, but equivalent, definitions of the index of ordinal variation have been described (Berry & Mielke, 1992a, 1992b; Blair & Lacy, 1996; Kvålseth, 1995a); here we follow Kvålseth's revised formulation (Kvålseth, 1995b).

entropy (Theil, 1972) and the interaction index (Reardon & Firebaugh, 2002), measures of diversity used in constructing the two-group information theory (*H*) and variance ratio segregation indices (James & Taeuber, 1985; Reardon & Firebaugh, 2002).

Two measures of ordinal segregation

Given the measures of ordinal variation, E_O and I_O, we define two measures of ordinal segregation from Equation (1) above:

$$H_O = 1 - \frac{\overline{E}_{Oj}}{E_O} = \sum_{j=1}^{J} \frac{t_j}{TE_O}\left(E_O - E_{Oj}\right) \tag{6}$$

and

$$R_O = 1 - \frac{\overline{I}_{Oj}}{I_O} = \sum_{j=1}^{J} \frac{t_j}{TI_O}\left(I_O - I_{Oj}\right) \tag{7}$$

where *j* indexes organizational units, t_j and T are the unit *j* and total population counts, respectively, and E_{Oj} and I_{Oj} are the ordinal entropy and ordinal variation in unit *j*. The first of these we name the *ordinal information theory index*, since it is an ordinal generalization of the categorical information theory index *H* (Theil, 1972), and is identical to *H* when K=2. The second we name the *ordinal variation ratio index*, since it is an ordinal generalization of the categorical variance ratio index (which goes by many names in the literature—see, e.g., James & Taeuber, 1985; Reardon & Firebaugh, 2002), and is identical to that index when K=2. Each of these indices is interpreted as the average difference in within-unit to overall ordinal variation, expressed as a ratio of the overall ordinal variation of the population. They differ in that H_O measures ordinal variation using the ordinal entropy E_O, while R_O uses the ordinal variation I_O.

Because both E_O and I_O are continuous, twice-differentiable, concave-down functions of the c_k in the domain $c_k \in [0,1]$,[5] both H_O and R_O range from a minimum of 0, obtained only when each unit j has the same distribution of x as the overall population (i.e., no segregation), to a maximum of 1, obtained only when there is no variation in x in any unit j, so that all the variation in x in the population lies between units (i.e., complete segregation) (Reardon & Firebaugh, 2002).

Ordinal segregation as an average of pairwise segregation

As derived above, H_O and R_O are similar to the variation ratio income segregation measures described in Section 1, though they rely only on the categorical counts to measure variation, and so do not require the estimation of the parameters of the income distribution. It is easy to show (see Appendix), however, that both H_O and R_O can be written as weighted averages of a set of K-1 pairwise segregation indices:

$$H_O = \frac{1}{(K-1)E_O}\sum_{k=1}^{K-1} E_k H_k \tag{8}$$

and

$$R_O = \frac{1}{(K-1)I_O}\sum_{k=1}^{K-1} I_k R_k \,, \tag{9}$$

---

[5] The second partial derivative of E_O with respect to any c_k (k=1,2,…,K-1) is $\frac{\partial^2 E_O}{\partial c_k^2} = \frac{-1}{c_k(1-c_k)} < 0$; the second partial derivative of I_O is $\frac{\partial^2 I_O}{\partial c_k^2} = -8$.

where the subscript k indicates variation or segregation computed between the two groups defined by the k^{th} threshold (i.e., H_k and R_k are segregation levels measured between a group consisting of all those with incomes in category k or below and a group consisting of all those in category k+1 or above). Since the ordinal measures E_O, I_O, H_O, and R_O are identical to their nominal counterparts when K=2, we drop the subscript "O" on these pairwise measures. These expressions indicate that the ordinal segregation measures can be seen as weighted averages of the binary segregation measures computed at each of the thresholds. We will return to a discussion of the interpretation of the weight terms later.

The fact that the two ordinal segregation indices can be written as weighted averages of a set of pairwise segregation indices enables us to better visualize what the indices measure (Figures 1 and 2). Figure 1 shows cumulative household income percentile density curves for each of the 176 census tracts in San Francisco County, CA (whose boundaries are identical to those of the city of San Francisco) in 2000.[6] Figure 2 shows the corresponding curves for the 613 tracts in Wayne County, MI (which includes Detroit). In both figures, the x-axis indicates both the local (i.e., San Francisco or Wayne County) income percentiles and the 15 income thresholds used in the 2000 census. Note that the income distribution in San Francisco is generally higher than in Wayne County—25% of households in San Francisco reported incomes greater than $100,000, compared to 12% of Wayne County households. If there were no income segregation in either county, each tract's cumulative household income percentile density curve would fall exactly on the 45-degree line (the heavy black line in each figure). If there were complete income segregation, each tract's curve would be a vertical line at some income level, indicating that within each tract all households have the same income. Thus, income segregation can be measured by the average deviation of the tract cumulative household income percentile density curves from their regional average (which is, by

[6] Data are obtained from Table P52 in SF3 from the 2000 Census. Although household income data are available at the block level, we aggregate to the tract level for these examples to reduce the number of curves shown in each figure.

definition, the 45-degree line). By this measure, Wayne County appears more segregated by income than San Francisco, since the variation of the tract cumulative density curves around the 45-degree line is greater in Wayne County.

Figures 1 and 2 here

The ordinal segregation indices H_O and R_O, when applied to income data, can be seen as measures of the variation of the tract cumulative household income percentile curves around the 45-degree line. At each income threshold reported in the Census (indicated by the vertical dashed lines in Figures 1 and 2), we know the value of each tract's cumulative household income density curve—that is, we know the percentage of households with incomes above and below each threshold. These data are used to compute the pairwise segregation between households with incomes above and below that threshold. These are then weighted (by E or I) and averaged across the thresholds to compute H_O and R_O.

Figures 3 and 4 show the pairwise household income segregation levels computed at each of the 15 Census 2000 thresholds for San Francisco and Wayne County, respectively. In addition, each figure illustrates the relative weight (dashed lines) that the pairwise segregation computed at each threshold is given in the calculation of the ordinal segregation measures (which are shown by the thin horizontal lines in each figure). First, note that segregation, as measured by either H or R, is relatively flat across most of the middle of the income percentile distribution in both places, but increases or decreases sharply at the extremes of the distribution, depending on which measure is used. Second, note that, as expected, measured segregation at each income percentile is generally higher in Wayne County than in San Francisco, regardless of which measure is used. Correspondingly, both ordinal segregation measures rank Wayne County as 16-17% more segregated than San Francisco: H_O=0.096 in San Francisco and H_O=0.113 in Wayne County; R_O=0.110 in San Francisco and R_O=0.128 in Wayne County. Finally, the shapes of the segregation curves differ slightly between San Francisco and Wayne County: in San Francisco, for example, segregation between those with income above and below the 25th percentile is higher than segregation between those with incomes above and below the 75th percentile, while the opposite is true in Wayne County.

Figures 3 and 4 here

The ordinal segregation measures for San Francisco and Wayne County shown in Figures 3 and 4 are not exactly comparable to one another. Because of the differences in the overall income distributions in the two counties, the Census-defined income thresholds do not fall at the same percentiles of the distributions. Thus, in San Francisco, the ordinal *H* and *R* are based more heavily on information about the segregation at thresholds in the 10th-50th percentile range (where 9 of 15 thresholds fall) than in the 50th-90th percentile range (where only 5 of 15 thresholds fall). In Wayne County, in contrast, three of the thresholds fall in the 95th-99th percentile range, and fewer at the low end. As a result, the ordinal segregation measured in San Francisco is not exactly comparable to that measured in Wayne County, because of differences in the underlying income distributions. Moreover, the measures clearly depend on the choice of thresholds—a different set of income thresholds would yield different measured levels of segregation. And finally, the measures are also clearly not invariant under changes in income that preserve the shape of the income distribution—a doubling of each household's income would have the effect of moving the thresholds to the left on the figure, while leaving the segregation curve (relative to income percentiles) unchanged, meaning that the computed ordinal segregation would depend much more on segregation levels at the low end of the percentile distribution.

The weights E_k and I_k in Equations (8) and (9) have substantively meaningful interpretations. The interpretation of E_k comes from information theory, which leads us to interpret E_k as the expected information content contained in H_k about a randomly-drawn individual from the population, where information is defined as $\ln(1/p_k)$, where p_k is the proportion of the population at or below the k^{th} threshold (Pielou, 1977; Theil, 1972). If p_k is near 0 or 1, then H_k contains little information about the segregation experienced by an individual, since it distinguishes among individuals only at one extreme of the income distribution. Conversely, if p_k is near 0.5, then H_k contains maximal information, since it distinguishes at the median of the distribution.

The weight I_k likewise has an appealing interpretation. For a given threshold k, the probability that two randomly-selected individuals from the population will have incomes on opposite sides of threshold k is $2p_k(1- p_k)$, which is proportional to I_k. Since the segregation level describes the extent of segregation between individuals on either side of the income threshold k, we can interpret Equation (9) as a weighted average of the segregation across each threshold, where there value at each threshold is weighted by how informative segregation measured at that threshold is for a randomly chosen pair of individuals.

Because both E_k and I_k have their maximum at p_k=0.5, and their minima at p_k=0 and p_k=1, H_O and R_O weight segregation between groups defined by the median of the income distribution most heavily, and segregation between the extreme income groups and the remainder least. Intuitively, this makes sense, since a segregation level computed between those above and below the 99th percentile, for example, tells us very little about the segregation between two randomly chosen individuals, while segregation between those above and below the median income tells us more about overall income segregation.

One implication of the above is that the segregation measures H_O and R_O can be seen both as variation ratio income segregation measures (since they are defined that way in Equations [6] and

[7]) and as categorical segregation measures (as defined in Equations [8] and [9]). Unfortunately, they share the flaws of the categorical measures, since Equations (8) and (9) make clear that they depend explicitly on the values of the K-1 income thresholds used to define the income categories (except in the special case where H_k and R_k are constant across all possible income thresholds, an unlikely scenario). Moreover, they share the flaws of the variation ratio measures, since the measures of ordinal variation on which they are based are not invariant under changes in the income distribution (unless the thresholds change as well so that they represent the same percentiles of the income distribution).

3. Rank-Order Measures of Income Segregation

The foregoing discussion illustrates that the ordinal segregation measures do not avoid the flaws of many existing approaches to measuring income segregation. They may improve on existing categorical measures, however, to the extent that they rely on multiple, relatively evenly-spaced thresholds, and because they weight the segregation at different thresholds more appropriately. Moreover, they may be useful measures when measuring segregation by some truly ordinal variable, where the thresholds have some substantive meaning (rather than a variable that is inherently continuous, but measured ordinally, like income). More importantly for our purposes, however, they provide the intuition for a related set of segregation measures that is free of their flaws.

Returning to the discussion of Figures 3 and 4, suppose we knew the value of the pairwise segregation computed at each point in the income distribution—that is, suppose we knew the shape of the function $H(p)$ or $R(p)$, where $H(p)$ and $R(p)$ are the pairwise segregation computed between those with incomes at or below the $100\times p^{th}$ income percentile and those above the $100\times p^{th}$

percentile.[7] Then we can compute the weighted average of the $H(p)$ and $R(p)$ over the interval (0,1) by extension from Equations (8) and (9). We define the *rank-order information theory index* (H_R):

$$\begin{aligned} H_R &= \int_0^1 \frac{E(p)}{\int_0^1 E(p)dp} H(p)dp \\ &= 2\ln(2)\int_0^1 E(p)H(p)dp \end{aligned} ; \qquad (10)$$

and the r*ank-order variation ratio index* (R_R):

$$\begin{aligned} R_R &= \int_0^1 \frac{I(p)}{\int_0^1 I(p)dp} R(p)dp \\ &= \frac{3}{2}\int_0^1 I(p)R(p)dp \end{aligned} . \qquad (11)$$

Intuitively, H_R and R_R are extensions of the ordinal segregation indices defined in Equations (6) and (7) to the case where we have an arbitrarily large number of categories—as many categories as individuals—in which case each individual's ranking operates as a distinct category.

Computing the Rank-Order Income Segregation Measures

To compute H_R and R_R, we must evaluate equations (10) and (11). The formulas for $E(p)$ and $I(p)$ are known by definition (Equations [4] and [5], where K=2). Thus, if we knew the functions

[7] We assume for simplicity that there are no ties on x—that is, no two observations have the same value of x. This ensures that the rank-ordering of x is unique and so segregation based on any given c is well-defined. In practice, however, this matters little when computing income segregation.

$H(p)$ or $R(p)$ on the interval (0,1), we could compute H_R and R_R without relying on an arbitrary set of thresholds. Of course, we do not know these functions, in general, but we may be able to estimate them from the H_k and R_k that we can measure.

We adopt the following strategy to estimate $H(p)$ or $R(p)$ on the interval (0,1). For each threshold k=1,2,…,K-1, we compute H_k or R_k and then plot them against the corresponding p_k, the cumulative proportions of the population with incomes equal to or below the threshold k (as in Figures 3 and 4) We then fit a polynomial of some order m to the observed points, using weighted least squares (WLS) regression and weighting each point by E_k^2 or I_k^2 (depending on whether we are fitting polynomial $H(p)$ or $R(p)$, respectively). Weighting by the square of the weight ensures that the fitted polynomial will fit best for p_k near 0.5, where H_k or R_k is weighted most, and that the weighted square root of the error variance is constant across p.

Suppose $H(p)$ and $R(p)$ are approximated by polynomial functions of order m:

$$\hat{H}(p) = \hat{\eta}_0 + \hat{\eta}_1 p + \hat{\eta}_2 p^2 + \cdots + \hat{\eta}_m p^m \tag{12}$$

and

$$\hat{R}(p) = \hat{\rho}_0 + \hat{\rho}_1 p + \hat{\rho}_2 p^2 + \cdots + \hat{\rho}_m p^m \,. \tag{13}$$

Substituting these into Equations (10) and (11), respectively, and evaluating the integrals, we get (see Appendix for derivations):

$$\hat{H}_R = \hat{\eta}_0 + \tfrac{1}{2}\hat{\eta}_1 + \tfrac{11}{36}\hat{\eta}_2 + \tfrac{5}{24}\hat{\eta}_3 + \ldots + \left[\frac{2}{(m+2)^2} + 2\sum_{n=0}^{m}\frac{(-1)^{m-n}\left({}_mC_n\right)}{(m-n+2)^2}\right]\hat{\eta}_m \tag{14}$$

and

$$\hat{R}_R = \hat{\rho}_0 + \tfrac{1}{2}\hat{\rho}_1 + \tfrac{3}{10}\hat{\rho}_2 + \tfrac{1}{5}\hat{\rho}_3 + \cdots + \tfrac{6}{(m+2)(m+3)}\hat{\rho}_m \,. \tag{15}$$

Thus, if we can estimate the parameters of polynomials approximating $H(p)$ and $R(p)$, we can easily estimate H_R and R_R as linear combinations of the coefficients of these polynomials. Moreover, to the extent that we can estimate $H(p)$ and $R(p)$ well on the interval (0,1) from the observed points, our estimate of segregation will not be biased by the choice of thresholds we have available. In general, we will not have information on the shape of $H(p)$ and $R(p)$ at the extreme ends of the income distribution (at points below and above the bottom and top thresholds), except via extrapolation. We can, however, assess the sensitivity of our estimates of H_R and R_R to alternative assumptions about the shape of $H(p)$ and $R(p)$ simply by assuming a range of possible shapes of the functional form. In general, the estimates of H_R and R_R will be relatively insensitive to assumptions about the shape of $H(p)$ and $R(p)$ at the ends of the income percentile distribution, since very little weight is given to $H(p)$ and $R(p)$ when p is near 0 or 1.[8]

Interpretation of the rank-order income segregation measures

We have defined two new income segregation indices, the *rank-order information theory index* (H_R) and the *rank-order variation ratio index* (R_R). These can each be interpreted in several ways. First, given an income-ranked population of N individuals (with no rank ties), we can define N-1 thresholds that each dichotomize the population into those with ranks above and below the given threshold. Now, the rank-order information theory index (H_R) is the weighted average of the N-1

[8] When weighting by E, the bottom and top deciles together carry only 7.5% of the total weight, for example; when weighting by I, the bottom and top deciles together carry only 5.6% of the total weight, so even if our estimates of $H(p)$ and $R(p)$were off by a large amount in the tails of the income percentile distribution, this error would contribute little error to the estimates of H_R and R_R.

values of the information theory index H obtained by computing the segregation between each of the pairs of groups defined by the thresholds, where the weight given to each value is proportional to the entropy corresponding to the percentile rank. Likewise, the rank-order variation ratio index (R_R) is the weighted average of the N-1 values of the variance ratio index R obtained by computing the segregation between each of the pairs of groups defined by the thresholds, where the weight given to each value is proportional to the probability that two randomly chosen individuals have ranks on either side of the corresponding threshold.

Second, the rank-order segregation indices can be interpreted as measures of the extent to which the cumulative income percentile density curves vary around their mean (the 45-degree line). The rank-order information theory index, for example, can be written as

$$H_R = 2\int_0^1 \left\{ \sum_{j=1}^{J} \frac{t_j}{T} \left[c_{jp} \ln\left(\frac{c_{jp}}{p}\right) + \left(1-c_{jp}\right) \ln\left(\frac{1-c_{jp}}{1-p}\right) \right] \right\} dp\,, \tag{16}$$

where c_{jp} is the cumulative proportion of those in neighborhood j with incomes at or below percentile $100\times p$. The term inside the brackets is akin to the Theil inequality measure, a measure of the deviation of c_{jp} from its mean (p). Summed over all neighborhoods j, this yields (inside the braces) a measure of the variation of the cumulative income percentile density functions at p. H_R is therefore a measure of the average variation of the neighborhood cumulative percentile density functions across percentiles.

Likewise, the rank-order relative diversity index can be written as

$$R_R = \frac{3}{2}\int_0^1 \left[\sum_{j=1}^{J} \frac{t_j}{T} \left(c_{jp} - p\right)^2 \right] dp\,. \tag{17}$$

Since the average of the c_{jp} across neighborhoods equals p, the term in the brackets is simply the variance of the c_{jp}'s at the point on the income distribution given by p. Thus, Equation (17) indicates that R_R can be interpreted as a measure of the average variance of the neighborhood cumulative percentile density functions across percentiles.

Third, the rank-order income segregation measures can be interpreted as variation ratios. Let E_R and I_R be the rank-order entropy and rank-order variation, respectively, when x is categorized into an arbitrarily large number of equal-sized categories (such as ranks). In the limit, as K gets arbitrarily large, we have $E_R=1/(2ln2)$ and $I_R=2/3$. Let E_R and I_R be the rank-order entropy and variation in neighborhood j, where the ranks are defined by the entire population. Thus we have

$$E_{Rj} = -\int_0^1 \left[c_{jp}\log_2(c_{jp}) + (1-c_{jp})\log_2(1-c_{jp})\right]dp \qquad (18)$$

and

$$I_{Rj} = \int_0^1 c_{jp}(1-c_{jp})dp\ . \qquad (19)$$

From Equations (10) and (18), we can then derive

$$H_R = \sum_{j=1}^{J}\frac{t_j}{T}\frac{\lfloor E_R - E_{Rj}\rfloor}{E_R}. \qquad (20)$$

Likewise, from Equations (11) and (19), we can derive

$$R_R = \sum_{j=1}^{J} \frac{t_j}{T} \frac{\lfloor I_R - I_{Rj} \rfloor}{I_R}. \tag{21}$$

Equations (20) and (21) show that the rank-order income segregation indices can be interpreted as variation ratio measures of segregation of the form given in Equation (1). The rank-order information theory index is interpreted as a measure of the difference between the average within-neighborhood rank-order entropy and that of the total population, expressed as a fraction of the rank-order entropy of the total population. Likewise, the rank-order variation ratio index is a measure of the ratio of within-neighborhood to overall rank-order variation. In other words, both indices measure how much less rank-order variation there is within neighborhoods than in the overall population.

Incorporating spatial proximity into measures of income segregation

The income segregation measures developed above do not take into account the spatial patterning of the organizational units (e.g. census tracts) in which income data are collected and reported. Moreover, they are inherently subject to the MAUP, since a different set of definitions of organizational boundaries may yield different computed levels of income segregation. However, because Equations (8) and (9) show that H_R and R_R are weighted averages of H and R computed across income percentiles, we can define explicitly spatial versions of H_R and R_R simply by substituting spatial versions of the pairwise H and R in Equations (8) and (9). Specifically, we define the *spatial rank-order information theory index* ($\widetilde{H}_R$):

$$\widetilde{H}_R = 2\ln(2)\int_0^1 E(p)\widetilde{H}(p)dp\,; \tag{22}$$

and the *spatial rank-order variation ratio index* ($\tilde{R}_R$):

$$\tilde{R}_R = \frac{3}{2}\int_0^1 I(p)\tilde{R}(p)dp, \tag{23}$$

where $\tilde{H}(p)$ and $\tilde{R}(p)$ are the pairwise spatial information theory index and the pairwise spatial relative diversity index (both defined in Reardon & O'Sullivan, 2004), respectively, computed between households with incomes above percentile 100×*p* and those with incomes at or below percentile 100×*p*. In practice, of course, the functions $\tilde{H}(p)$ and $\tilde{R}(p)$ can be estimated from the values of $\tilde{H}_k$ and $\tilde{R}_k$ computed at each of the *K*-1 income thresholds used to define income categories, just as we propose in the aspatial case above. Thus, the spatial versions of the rank-order segregation indices introduce no additional complexity to the measures, other than the additional computational burden of computing spatial versions of $\tilde{H}_k$ and $\tilde{R}_k$.

Both of the rank-order income segregation measures satisfy the three general criteria we desire. First, because they are based on the rank ordering of the income distribution, rather than actual values of income, they are insensitive to a rank-order preserving changes in incomes. Second, even if income data are reported in categories, the measures can be estimated with considerable precision, so long as sufficient income categories are reported. Thus, even though the reported income categories may change from Census to Census, or may vary from country to country, the rank-order measures are unaffected by these changes, both in principle, and in practice. Third, the rank-order segregation measures are easily adapted to incorporate a flexible spatial proximity function, since the two-group segregation measures from which they are computed can be made to incorporate spatial proximity (Reardon et al., 2006; Reardon & O'Sullivan, 2004).

The rank-order income segregation measures H_R and R_R (and their spatial counterparts) share a number of properties, which is not surprising, given that both are derived from similar approaches, and differ only in the variation functions they use. In addition to satisfying our three general criteria, both also inherit many of the properties of the pairwise categorical segregation indices from which they are derived (see James & Taeuber, 1985; Reardon & Firebaugh, 2002; Reardon & O'Sullivan, 2004). Both, for example, are additively spatially decomposable. Neither satisfies the composition invariance criterion advocated by James and Taueber (1985), however, but it is not clear that this criterion has any meaning in the case of rank-order income segregation.[9] The spatial pairwise information theory index, however, satisfies a stronger form of the exchange criterion than does the corresponding relative diversity index (Reardon & O'Sullivan, 2004). Other than this criterion, we find little to recommend one index over the other, and, in fact, we shall see that they yield substantively very similar results.

4. Empirical Examples

We illustrate the properties of the two rank-order segregation measures using household income data from several metropolitan areas in the U.S. in 2000—Atlanta, Denver, Minneapolis, New York, Pittsburgh, and San Jose—chosen to illustrate a diverse range of metropolitan areas.[10] In the empirical illustrations that follow, we demonstrate several features of the proposed measures. First, we show that the polynomial approximations to H_R and R_R yield very precise estimates of these quantities. Second, we show how a display of the income segregation threshold at different

[9] The usual composition invariance criterion states that if the number of members of one group is multiplied by a constant in each neighborhood, segregation is unchanged. In the case of rank order segregation, it is, in general, mathematically impossible to change neighborhood compositions in a way that meets the composition invariance criterion.

[10] The results reported here rely on counts of households in each of 16 income categories by census block group (the smallest level of aggregation for which household income data are reported in the 2000 Census), obtained from Summary File 3, Table 52, of the 2000 U.S. Census. Metropolitan area boundaries are based on definitions published by OMB in 2003.

spatial scales can reveal considerable insight regarding the spatial patterning of household income in a region. In the results that follow, we illustrate patterns of income segregation as measured by H_R, since there are slight reasons to prefer H_R over R_R (discussed above), though we show corresponding results for R_R in additional tables. None of the substantive results change if we use R_R rather than H_R.

Table 1 reports estimated household income segregation levels for the 6 metropolitan areas. The first column in each panel of the table reports the ordinal income segregation measures H_O and R_O, while the subsequent columns report the rank-order income segregation measures estimated based on polynomial approximations of orders M=2 through M=10. First, note that the rank-order measures are remarkably stable, regardless of the order of polynomial used. This is largely because the functions $H(p)$ and $R(p)$ are relatively smooth functions, well-approximated by low-order polynomials, so the higher-order terms add little to the fit of the curves. The estimated income segregation levels vary the most with the order of the polynomial in the case of the New York metropolitan area, where the values of H_R vary from a high of 0.1478 (when M=2) to a low of 0.1462 (when M=3)—a difference of slightly more than one percent. Figure 5 illustrates the values of H_k at each of the 15 thresholds for the New York metropolitan area, as well as the fitted polynomials of order M=2, 3,…, 10. Note that for polynomials of order 4 or higher, the curves fit the points extremely well through most of the range. At the extremes of the income distribution, where there is no information to fit the curves, polynomials of different orders produce substantially different curves, but this variation has almost no effect on the estimates of H_R, since $H(p)$ is weighted very little in these regions.

Table 1 about here

Figure 5 about here

Second, Table 1 indicates that the ordinal measures H_O and R_O are often a reasonably good estimate of the rank-order measures H_R and R_R, though not in all cases. In San Jose, for example, R_O is 10 percent lower than the values of R_R. In Pittsburgh, in contrast, H_O is 3% larger than the values of H_R. Moreover, Atlanta appears more segregated than Minneapolis on the basis of H_O, but less segregated on the basis of H_R (when estimated with any order polynomial greater than 2). Given that the error in the rank order measures due to the polynomial approximation is generally less than one percent of the true value, the error in the ordinal measures (due to uneven spacing of the thresholds) can be considerable.[11]

We next illustrate the evidence provided by the spatial rank order information theory index ($\tilde{H}_R$). Following the example of Reardon and colleagues (Lee et al., 2006; Reardon et al., 2006; Reardon & O'Sullivan, 2004), we compute $\tilde{H}_R$ using a biweight kernel proximity function with varying radii ranging from 500m (the scale of a pedestrian neighborhood) to 4000m (a radius that encompasses most regular activities, such as shopping, attending high school, or attending religious services; see, e.g., Sastry, Pebley, & Zonta, 2002).[12] We can interpret $\tilde{H}_R$ at a 500m radius, for example, as indicating how much less household income variation there is in the 500m-radius local environment of the average household than there is in the metropolitan region as a whole. Like Reardon and colleagues, we also compute a measure of the granularity, or scale, of residential income segregation by computing the ratio of $\tilde{H}_R$ at a 4000m radius to $\tilde{H}_R$ at a 500m radius. This

---

[11] In additional analyses (not shown), we examined the stability of the estimates of H_R and R_R over polynomials of order 2 to 10 for the 100 largest metropolitan areas in 2000. In general, we found that the rank order segregation indices changed very little as we increased the order of the polynomial. Above order M=6, increasing the order of the approximating polynomial yielded an average absolute change of ±0.0002 in H_R, for example (a change of less than two-tenths of a percent in H), while changes in R_R were consistently smaller. Consequently, in the remainder of our empirical examples, we estimate H_R and R_R using eighth-order polynomial approximations to $H(p)$ and $R(p)$, though we get substantively indistinguishable results using polynomials of any order $M \geq 2$.

[12] The spatial segregation measures reported here—including estimation of the population densities and computation of segregation levels—are computed using a macro we have written in Visual Basic for Applications (VBA) and run within ArcGIS 9.1 software (Environmental Systems Research Institute, 2005).

ratio measure indicates the extent to which the segregation of micro (500m radius) environments is due to larger scale macro-segregation (Lee et al., 2006; Reardon et al., 2006).

Table 2 about here

Table 2 reports values, for each of the 6 metropolitan areas, of the spatial rank order information theory index and the spatial relative diversity index, each computed at radii of 500m, 1000m, 2000m, and 4000m. In addition, we include the corresponding aspatial measures, computed using block groups as neighborhood units. Finally, the last column indicates the granularity ratio for each metropolitan area. Several facts are notable in the table. First, income segregation declines with scale, though at different rates across metropolitan areas. In Atlanta, income segregation computed using 4000m-radius local environments is two-thirds (H_R ratio=0.65; R_R ratio=0.66) of segregation computed using 500m-radius environments; in San Jose, however, the corresponding ratios are H_R=0.42 or R_R=0.43. Likewise, while Denver reveals more income segregation than Atlanta at a 500m radius, it is less segregated than Atlanta at 2000m or 4000m radius, an illustration of the scale-sensitive nature of segregation (Reardon et al., 2006).

Second, segregation computed among block groups is generally lower than segregation at a 500m radius—since block groups, on average, are larger than a 500m radius environment. In New York, however, block group segregation is higher than 500m-radius spatial segregation, a result of the high population density and small census block groups in the New York region. We note further that, relying on block group segregation measures would lead us to conclude that New York ranks second in income segregation among the 6 metropolitan areas here, while the spatial measures would rank it fourth of 6. Finally, we note that substantive patterns and comparisons in Table 2 are virtually identical regardless of whether we base them on H_R or R_R.

The results in Table 2 provide summary measures of spatial income segregation. However, we can examine the patterns of income segregation in more detail for each of the metropolitan areas in Figures 6-11. In each figure, we plot the function $\tilde{H}(p)$—the pairwise spatial information theory index computed at the threshold defined by income percentile *p*—estimated at the four radii (solid lines). In addition, we plot the corresponding function $H(p)$ based on census block groups (dashed line). Finally, we plot the granularity ratio as a function of *p* (dotted line), in order to examine how the granularity of segregation varies across the range of income percentiles. We include figures for each of the 6 metropolitan areas, though we do not discuss each in detail.

Figure 6 illustrates income segregation patterns in the Atlanta metropolitan area. First, note that segregation between the highest-earning households and others is higher than segregation between the lowest-earning households and other. At a 500m radius, for example, segregation between the top decile of households and all others is about 0.20, while segregation between the bottom decile and others is about 0.14.[13] Second, note that the granularity ratio is quite high, and very stable across the range of income percentiles. This indicates that the majority of income segregation in the Atlanta metropolitan area, regardless of what income threshold we use, is due to macro-segregation patterns—large-scale differences in the spatial distribution of households by income across the region.

Figure 9 (New York) illustrates some similar patterns as evident in Atlanta. The segregation gradient at the high end of the income distribution, however, is quite steep, however (see also Pittsburgh, Figure 10), indicating that the highest-income households are substantially more segregated from other households than are the lowest-income households. Likewise, in the San Jose

[13] Note that we can only validly compare segregation at percentile *p* to segregation at percentile 100-*p*, since a valid comparison between other thresholds (say, between segregation at *p=10* and at *p*=50) requires a measure of segregation that is "composition invariant" in some meaningful sense. The usual definitions of composition invariance (Coleman, Hoffer, & Kilgore, 1982; James & Taeuber, 1985; Reardon & O'Sullivan, 2004), however, do not apply to the rank order income segregation measures. A fuller discussion of this issue is beyond the scope of this paper; in the absence of such a discussion, we limit our discussion here to comparisons that are meaningful in the absence of composition invariance.

metropolitan area (which includes much of Silicon Valley), for example (Figure 11), although overall segregation levels are lower than in New York, segregation of higher-income households from others is considerably higher than segregation of lower-income households from others. San Jose also demonstrates much more variation in the granularity of segregation across the range of incomes—segregation of high-income households is largely accounted for by macro-scale segregation (meaning there are large spatial regions, in places like Palo Alto, where high-income households are concentrated), while segregation of lower-income households is much more localized (largely in the city of San Jose).

5. Conclusion

In this paper we have proposed two measures of income segregation, the *rank-order information theory index* (H_R), and the *rank-order variation ratio index* (H_R), as well as their spatially-explicit counterparts. These indices have several appealing features. First, they are relatively easy to compute, since they require (in the aspatial case) simply computing a series of pairwise segregation values using existing measures of segregation (H and R), fitting a polynomial regression line using WLS, and then computing a linear combination of the estimated parameters. Second, the measures are easily adapted to account for spatial proximity, following the approach of Reardon and O'Sullivan (2004). In the spatial case, the computational steps are the same, but the pairwise segregation indices must be computed using some spatially-sensitive method. Third, the measures are largely insensitive to the set of income thresholds that define the categories in which income is reported, so long as the thresholds span most of the range of income percentiles. As a result, they do not require us to make assumptions about the shape of income distributions. Finally, the measures are insensitive to rank-preserving changes in income, since the measures are based on the ranks of incomes rather than their numerical values.

Appendix

Derivation of Equations (8) and (9)

Let E_k indicate the entropy of a population that is divided into two groups, where one group consists of all those in income category k or below, and the other groups consists of all those in categories k+1 or above:

$$E_k = -c_k \log_2(c_k) - (1-c_k)\log_2(1-c_k).$$

Likewise, let H_k denote the information theory index (Theil, 1972) computed between these two groups. Then we can rewrite Equation (6) as:

$$\begin{aligned}
H_O &= \sum_{j=1}^{J} \frac{t_j}{TE_O}\left(E_O - E_{Oj}\right) \\
&= \sum_{j=1}^{J} \frac{t_j}{TE_O}\left[\frac{1}{K-1}\sum_{k=1}^{K\text{-}1}\left(E_k - E_{kj}\right)\right] \\
&= \sum_{k=1}^{K-1} \frac{E_k}{(K-1)E_O}\left[\sum_{j=1}^{J}\frac{t_j}{TE_k}\left(E_k - E_{kj}\right)\right] . \\
&= \sum_{k=1}^{K-1} \frac{E_k}{(K-1)E_O} H_k
\end{aligned}$$

Note that Equation (4) implies that

$$\sum_{k=1}^{K-1} E_k = (K-1)E_O ,$$

so the above shows that H_O can be expressed as a weighted average of the K-1 pairwise H_k's. An identical derivation, with I_k and R_k substituted for E_k and H_k, yields Equation (9).

Derivation of Equation (14)

Substituting Equation (12) into (10), rearranging terms, and integrating, yields

$$
\begin{aligned}
\hat{H}_R &= 2\ln(2)\int_0^1 E(p)\hat{H}(p)dp \\
&= -2\int_0^1 \left[p\ln p + (1-p)\ln(1-p)\right]\left[\sum_{m=0}^{M}\hat{\eta}_m p^m\right]dp \\
&= -\sum_{m=0}^{M} 2\hat{\eta}_m\left[\int_0^1 p^{m+1}\ln p dp + \int_0^1 p^m(1-p)\ln(1-p)dp\right] \\
&= -\sum_{m=0}^{M} 2\hat{\eta}_m\left[\int_0^1 p^{m+1}\ln p dp + \int_0^1 [1-(1-p)]^m(1-p)\ln(1-p)dp\right] \\
&= -\sum_{m=0}^{M} 2\hat{\eta}_m\left[p^{m+1}\ln p dp + \int_0^1\left[\sum_{n=0}^{m}(-1)^{m-n}\left({}_mC_n\right)(1-p)^{m-n+1}\right]\ln(1-p)dp\right] \\
&= -\sum_{m=0}^{M} 2\hat{\eta}_m\left[p^{m+1}\ln p dp + \sum_{n=0}^{m}(-1)^{m-n}\left({}_mC_n\right)\int_0^1 (1-p)^{m-n+1}\ln(1-p)dp\right] \\
&= -\sum_{m=0}^{M} 2\hat{\eta}_m\left[\begin{array}{l} p^{m+2}\left(\dfrac{\ln p}{m+2}-\dfrac{1}{(m+2)^2}\right)\Bigg|_0^1 \\ +\sum_{n=0}^{m}(-1)^{m-n}\left({}_mC_n\right)\left[-(1-p)^{m-n+2}\left(\dfrac{\ln(1-p)}{m-n+2}-\dfrac{1}{(m-n+2)^2}\right)\Bigg|_0^1\right]\end{array}\right]. \\
&= \sum_{m=0}^{M} 2\hat{\eta}_m\left[\frac{1}{(m+2)^2}+\sum_{n=0}^{m}\frac{(-1)^{m-n}\left({}_mC_n\right)}{(m-n+2)^2}\right]
\end{aligned}
$$

So, for values of M, we have:

$$
\begin{aligned}
M &= 0 \Rightarrow \hat{H}_R = \hat{\eta}_0 \\
M &= 1 \Rightarrow \hat{H}_R = \hat{\eta}_0 + \tfrac{1}{2}\hat{\eta}_1 \\
M &= 2 \Rightarrow \hat{H}_R = \hat{\eta}_0 + \tfrac{1}{2}\hat{\eta}_1 + \tfrac{11}{36}\hat{\eta}_2 \\
M &= 3 \Rightarrow \hat{H}_R = \hat{\eta}_0 + \tfrac{1}{2}\hat{\eta}_1 + \tfrac{11}{36}\hat{\eta}_2 + \tfrac{5}{24}\hat{\eta}_3 \\
M &= 4 \Rightarrow \hat{H}_R = \hat{\eta}_0 + \tfrac{1}{2}\hat{\eta}_1 + \tfrac{11}{36}\hat{\eta}_2 + \tfrac{5}{24}\hat{\eta}_3 + \tfrac{137}{900}\hat{\eta}_4 \\
&\vdots \\
M &= m \Rightarrow \hat{H}_R = \hat{\eta}_0 + \tfrac{1}{2}\hat{\eta}_1 + \tfrac{11}{36}\hat{\eta}_2 + \ldots + \left[\frac{2}{(m+2)^2} + 2\sum_{n=0}^{m}\frac{(-1)^{m-n}\left({}_mC_n\right)}{(m-n+2)^2}\right]\hat{\eta}_m
\end{aligned}
$$

Derivation of Equation (15)

Substituting Equation (13) into (11), rearranging terms, and integrating, yields

$$
\begin{aligned}
\hat{R}_R &= \frac{3}{2}\int_0^1 I(p)\hat{R}(p)dp \\
&= 6\int_0^1 \left[p(1-p)\right]\left[\sum_{m=0}^{M}\hat{\rho}_m p^m\right]dp \\
&= \sum_{m=0}^{M} 6\hat{\rho}_m\left[\int_0^1 p^{m+1}(1-p)dp\right] \\
&= \sum_{m=0}^{M} 6\hat{\rho}_m\left[\left(\frac{p^{m+2}}{m+2}-\frac{p^{m+3}}{m+3}\right)\Bigg|_0^1\right] \\
&= \sum_{m=0}^{M} \frac{6}{(m+2)(m+3)}\hat{\rho}_m
\end{aligned}
\quad .
$$

So, for values of *M*, we have:

$$
\begin{aligned}
&M = 0 \Rightarrow \hat{R}_R = \hat{\rho}_0 \\
&M = 1 \Rightarrow \hat{R}_R = \hat{\rho}_0 + \tfrac{1}{2}\hat{\rho}_1 \\
&M = 2 \Rightarrow \hat{R}_R = \hat{\rho}_0 + \tfrac{1}{2}\hat{\rho}_1 + \tfrac{3}{10}\hat{\rho}_2 \\
&M = 3 \Rightarrow \hat{R}_R = \hat{\rho}_0 + \tfrac{1}{2}\hat{\rho}_1 + \tfrac{3}{10}\hat{\rho}_2 + \tfrac{1}{5}\hat{\rho}_3 \\
&M = 4 \Rightarrow \hat{R}_R = \hat{\rho}_0 + \tfrac{1}{2}\hat{\rho}_1 + \tfrac{3}{10}\hat{\rho}_2 + \tfrac{1}{5}\hat{\rho}_3 + \tfrac{1}{7}\hat{\rho}_4 \\
&\vdots \\
&M = m \Rightarrow \hat{R}_R = \hat{\rho}_0 + \tfrac{1}{2}\hat{\rho}_1 + \tfrac{3}{10}\hat{\rho}_2 + \ldots + \frac{6}{(m+2)(m+3)}\hat{\rho}_m
\end{aligned}
\quad .
$$

References

Berry, K. J., & Mielke, P. W., Jr. (1992a). Assessment of variation in ordinal data. *Perceptual and Motor Skills, 74*, 63-66.

Berry, K. J., & Mielke, P. W., Jr. (1992b). Indices of ordinal variation. *Perceptual and Motor Skills, 74*, 576-578.

Blair, J., & Lacy, M. G. (1996). Measures of variation for ordinal data as functions of the cumulative distribution. *Perceptual and Motor Skills, 82*, 411-418.

Chakravorty, S. (1996). A measurement of spatial disparity: The case of income inequality. *Urban Studies, 33*(9), 1671-1686.

Coleman, J., Hoffer, T., & Kilgore, S. (1982). Achievement and segregation in secondary schools: A further look at public and private school differences. *Sociology of Education, 55*((April/July)), 162-182.

Coulton, C. J., Chow, J., Wang, E. C., & Su, M. (1996). Geographic concentration of affluence and poverty in 100 metropolitan areas, 1990. *Urban Affairs Review, 32*, 186-216.

Davidoff, T. (2005). *Income sorting: Measurement and decomposition* (Working paper). Berkeley, CA: UC-Berkeley.

Dawkins, C. J. (forthcoming). Space and the measurement of income segregation. *Journal of Regional Science*.

Environmental Systems Research Institute. (2005). ArcGIS: Release 9.1. Redlands, CA: Environmental Systems Research Institute.

Farley, R. (1977). Residential segregation in urbanized areas of the United States in 1970: An analysis of social class and racial differences. *Demography, 14*(4), 497-518.

Fong, E., & Shibuya, Kumiko. (2000). The spatial separation of the poor in Canadian cities. *Demography, 37*(4), 449-459.

Hardman, A., & Ioannides, Y. M. (2004). Neighbors' income distribution: Economic segregation and mixing in U.S. urban neighborhoods. *Journal of Housing Economics, 13*, 368-382.

Ioannides, Y. M. (2004). Neighborhood income distributions. *Journal of Urban Economics, 56*(3), 435-457.

Ioannides, Y. M., & Seslen, T. N. (2002). Neighborhood wealth distributions. *Economics Letters, 76*, 357-367.

James, D. R., & Taeuber, K. E. (1985). Measures of segregation. *Sociological Methodology, 14*, 1-32.

Jargowsky, P. A. (1996). Take the money and run: Economic segregation in U.S. metropolitan areas. *American Sociological Review, 61*(6), 984-998.

Jargowsky, P. A. (1997). *Poverty and place: Ghettos, barrios, and the American city*. New York: Russell Sage Foundation.

Jargowsky, P. A., & Kim, J. (2005). *A measure of spatial segregation: The generalized neighborhood sorting index* (NPC Working paper No. 05-3). Ann Arbor, MI: National Poverty Center, University of Michigan.

Jenkins, S. P., Micklewright, J., & Schnepf, S. V. (2006). *Social segregation in schools: An international comparison.* Paper presented at the 29th General Conference of The International Association for Research in Income and Wealth, Joensuu, Finland.

Kvålseth. (1995a). Coefficients of variation for nominal and ordinal categorical data. *Perceptual and Motor Skills, 80*, 843-847.

Kvålseth. (1995b). Comment on the coefficient of ordinal variation. *Perceptual and Motor Skills, 81*, 621-622.

Lee, B. A., Reardon, S. F., Firebaugh, G., Farrell, C. R., Matthews, S. A., & O'Sullivan, D. (2006, March 30-April 1). *Beyond the Census Tract: Patterns and Determinants of Racial Residential Segregation at Multiple Scales.* Paper presented at the Annual Meeting of the Population

Association of America, Los Angeles, CA.

Massey, D. S. (1996). The age of extremes: Concentrated affluence and poverty in the twenty-first century. *Demography, 33*(4), 395-412.

Massey, D. S., & Eggers, M. L. (1990). The ecology of inequality: Minorities and the concentration of poverty, 1970-1980. *American Journal of Sociology, 95*(5), 1153-1188.

Massey, D. S., & Eggers, M. L. (1993). The spatial concentration of affluence and poverty during the 1970s. *Urban Affairs Quarterly, 29*, 299-315.

Massey, D. S., & Fisher, M. J. (2003). The geography of inequality in the United States, 1950-2000. *Brookings-Wharton Papers on Urban Affairs*(4), 1-40.

Meng, G., Hall, G. B., & Roberts, S. (2006). Multi-group segregation indices for measuring ordinal classes. *Computers, Environment and Urban Systems, 30*(3), 275-299.

Openshaw, S. (1984). *The Modifiable Area Unit Problem* (Vol. 38). Norwich: Geo books.

Pendall, R., & Carruthers, J. I. (2003). Does density exacerbate income segregation? Evidence from U.S. metropolitan areas, 1980 to 2000. *Housing Policy Debate, 14*(4), 541-589.

Pielou, E. C. (1977). *Mathematical ecology* (2nd ed.). New York: John Wiley & Sons.

Reardon, S. F., & Firebaugh, G. (2002). Measures of multi-group segregation. *Sociological Methodology, 32*, 33-67.

Reardon, S. F., Matthews, S. A., O'Sullivan, D., Lee, B. A., Firebaugh, G., & Farrell, C. R. (2006). *The Segregation Profile: Investigating How Metropolitan Racial Segregation Varies by Spatial Scale.* Paper presented at the Annual Meeting of the Population Association of America, Los Angeles, CA.

Reardon, S. F., & O'Sullivan, D. (2004). Measures of spatial segregation. *Sociological Methodology, 34*, 121-162.

Sastry, N., Pebley, A. R., & Zonta, M. (2002). *Neighborhood definitions and the spatial dimension of daily life* *in Los Angeles* (No. ccpr-033-04): California Center for Population Research, On-Line Working Paper Series.

Telles, E. E. (1995). Structural sources of socioeconomic segregation in Brazilian metropolitan areas. *American Journal of Sociology, 100*(5), 1199-1223.

Theil, H. (1972). *Statistical decomposition analysis* (Vol. 14). Amsterdam: North-Holland Publishing Company.

Waitzman, N. J., & Smith, K. R. (1998). Separate but lethal: The effects of economic segregation on mortality in metropolitan America. *The Milbank Quarterly, 76*(3), 341-373.

Watson, T. (2006). *Metropolitan growth, inequality, and neighborhood segregation by income* (Working paper). Williamstown, MA: Williams College.

Wheeler, C. H. (2006). *Urban decentralization and income inequality: Is sprawl associated with rising income segregation across neighborhoods?* (FRB of St. Louis Working paper No. 2006-037A). St. Louis, MO: Federal Reserve Bank of St. Louis.

Wheeler, C. H., & La Jeunesse, E. A. (2006). *Neighborhood income inequality* (FRB of St. Louis Working paper No. 2006-039A). St. Louis, MO: Federal Reserve Bank of St. Louis.

Figure 1

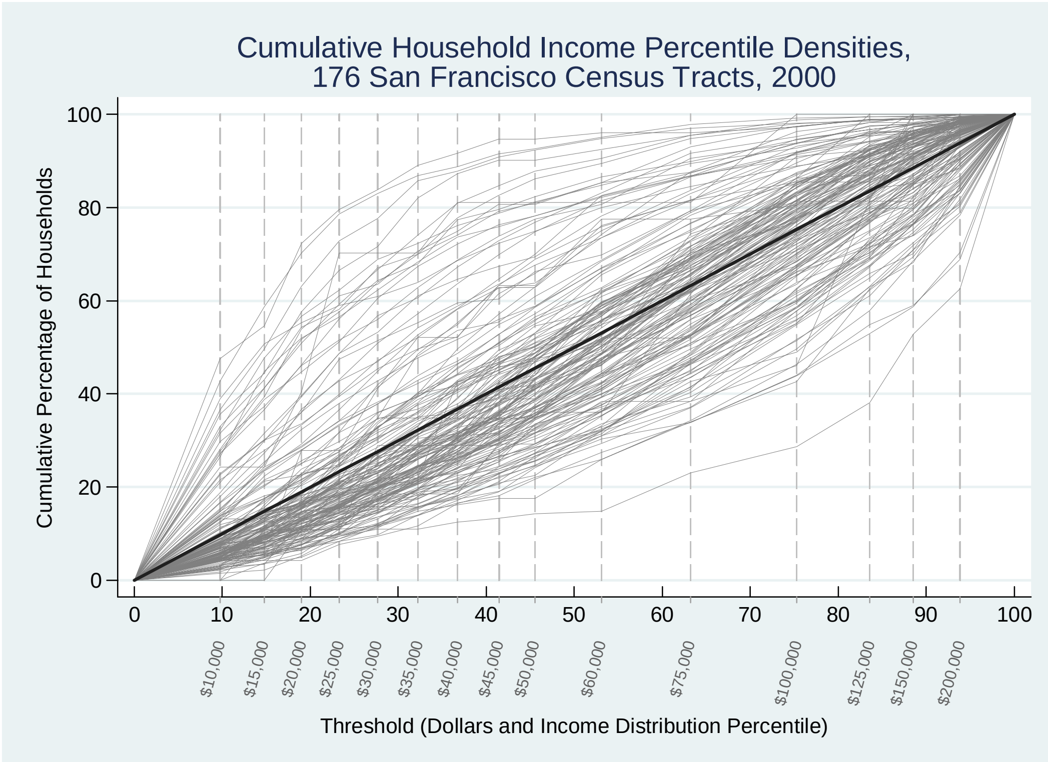

Figure 2

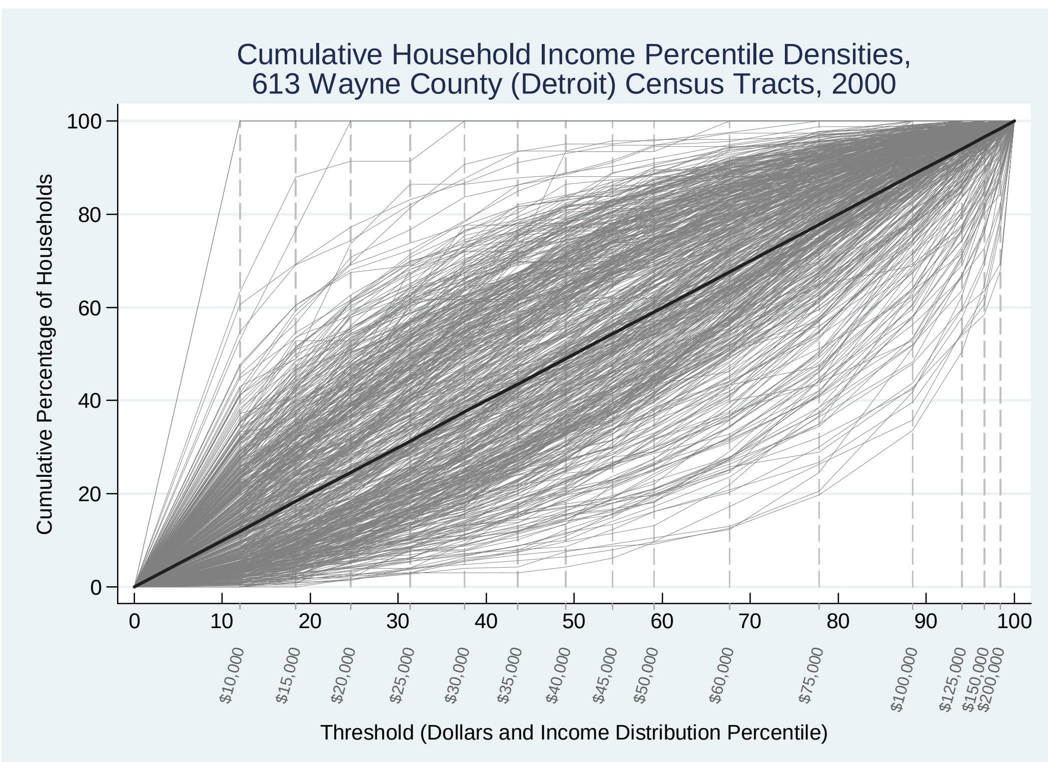

Figure 3

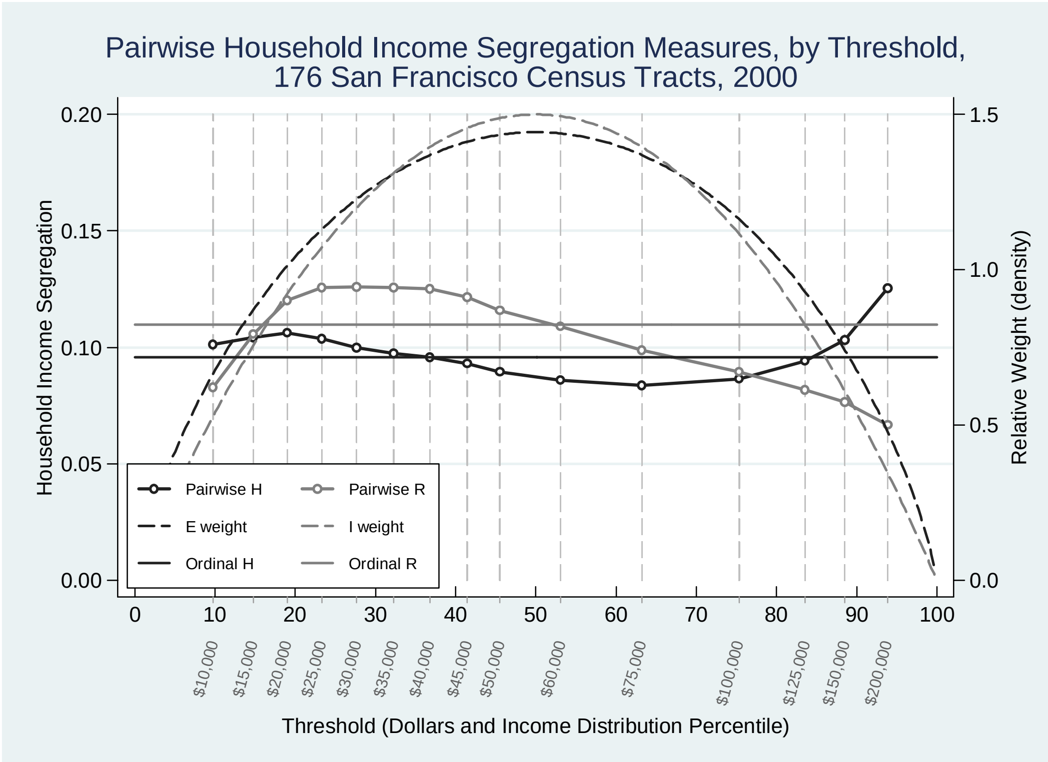

Figure 4

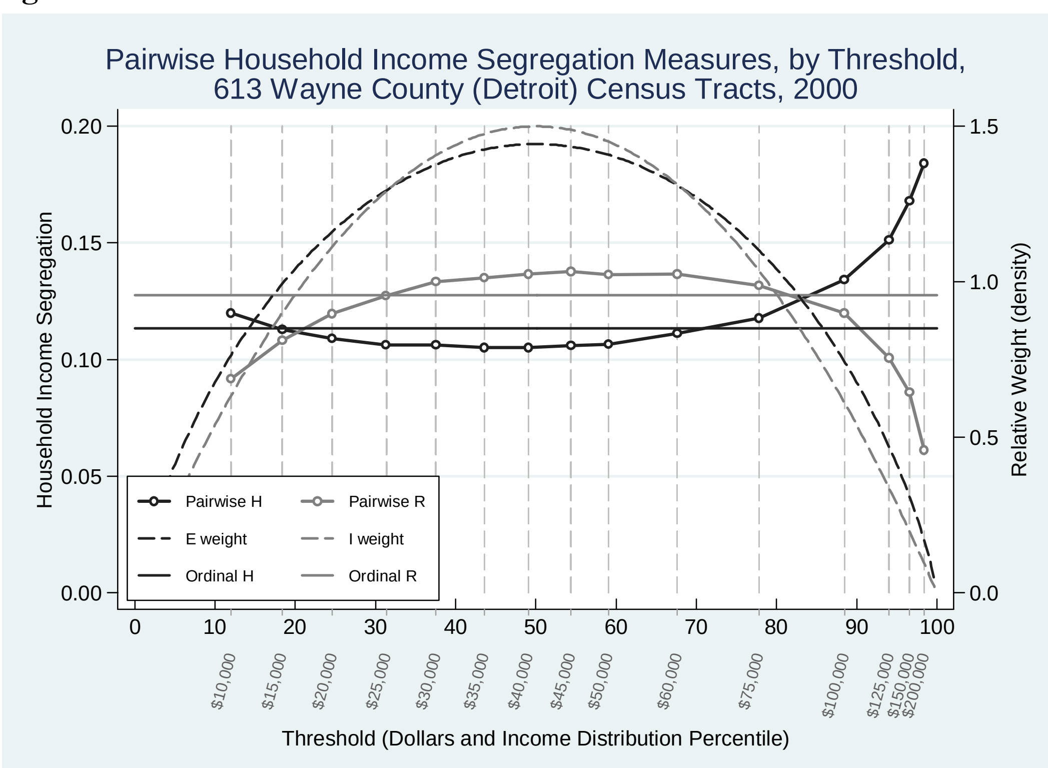

Figure 5

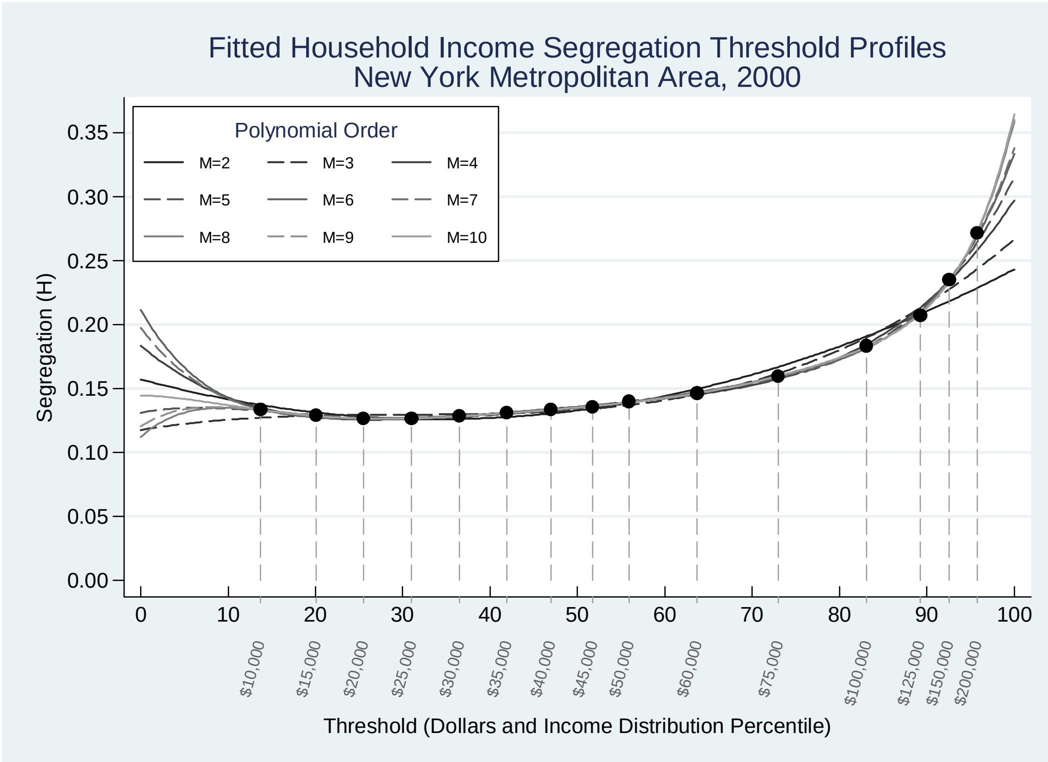

Fitted Household Income Segregation Threshold Profiles
New York Metropolitan Area, 2000
Polynomial Order
M=2
M=3
M=4
M=5
M=6
M=7
M=8
M=9
M=10
Segregation (H)
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0
10
20
30
40
50
60
70
80
90
100
$10,000
$15,000
$20,000
$25,000
$30,000
$35,000
$40,000
$45,000
$50,000
$60,000
$75,000
$100,000
$125,000
$150,000
$200,000
Threshold (Dollars and Income Distribution Percentile)

Figure 6

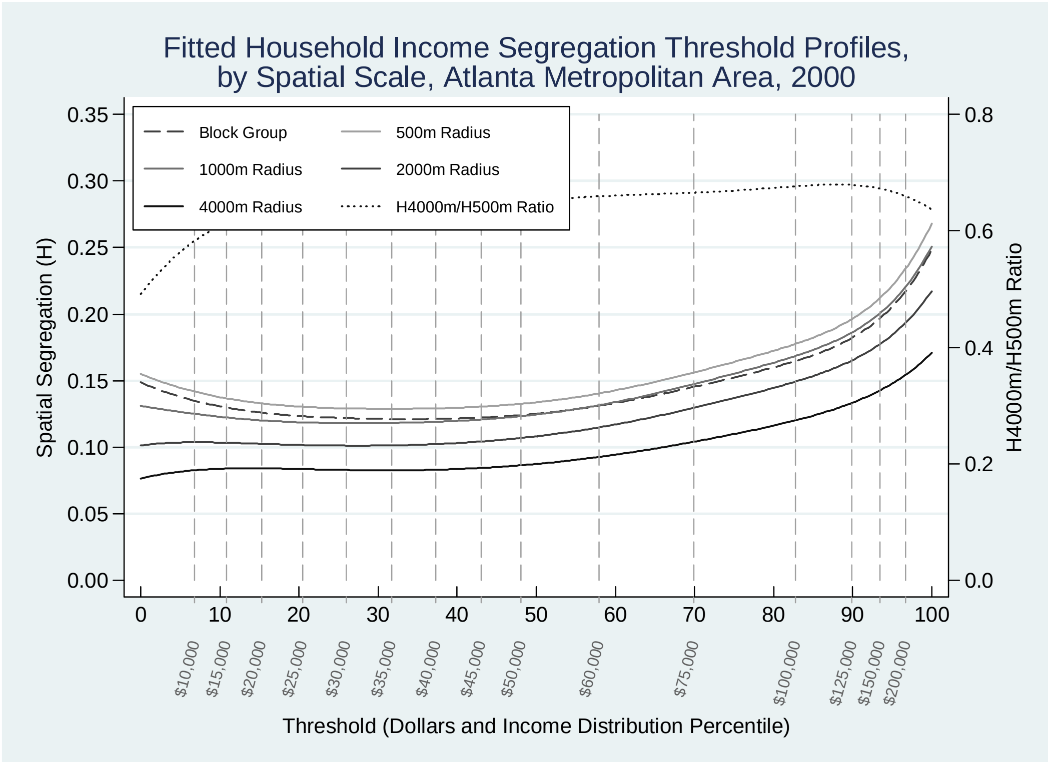

Figure 7

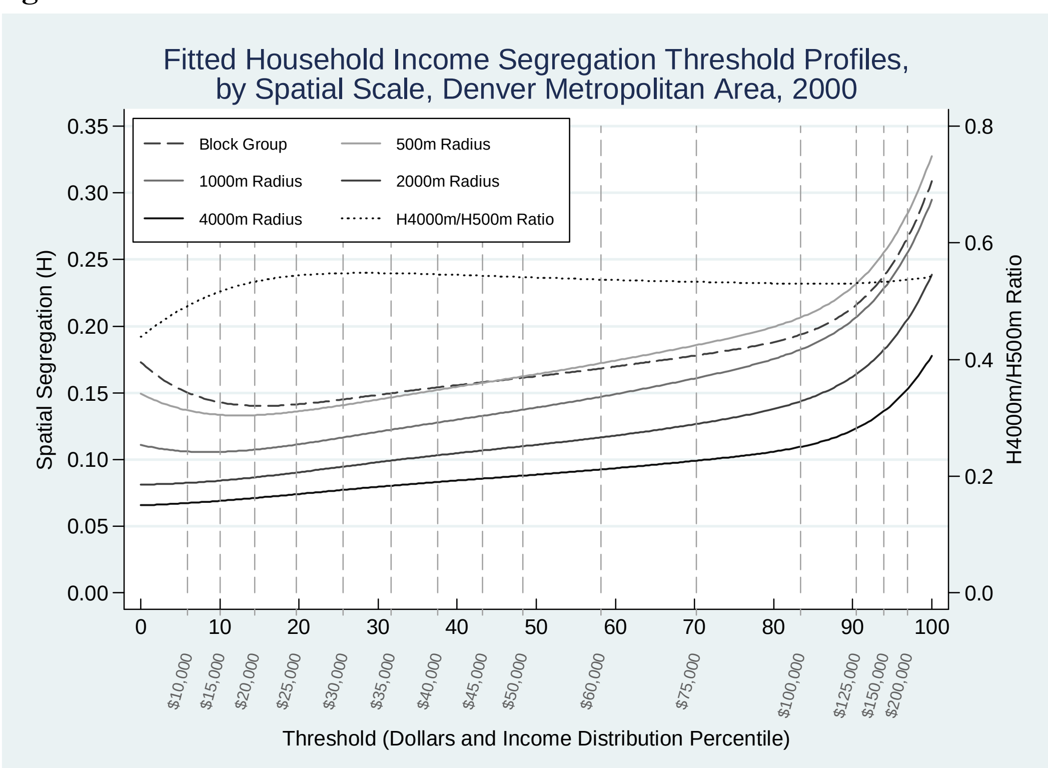

Figure 8

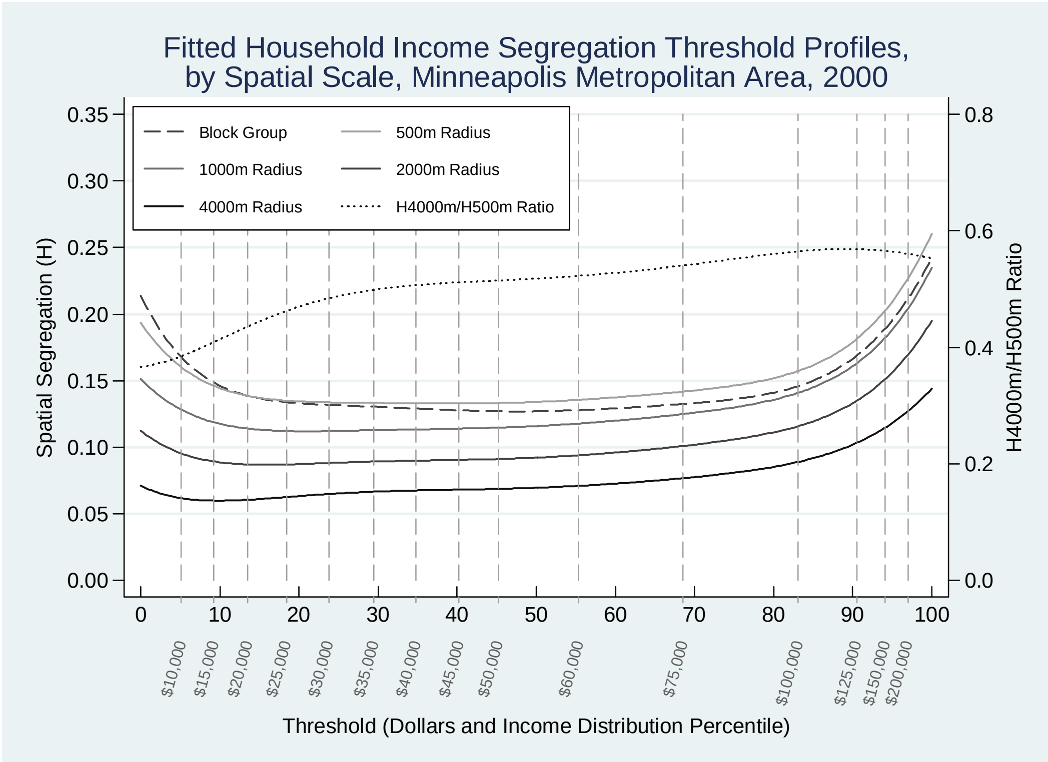

Figure 9

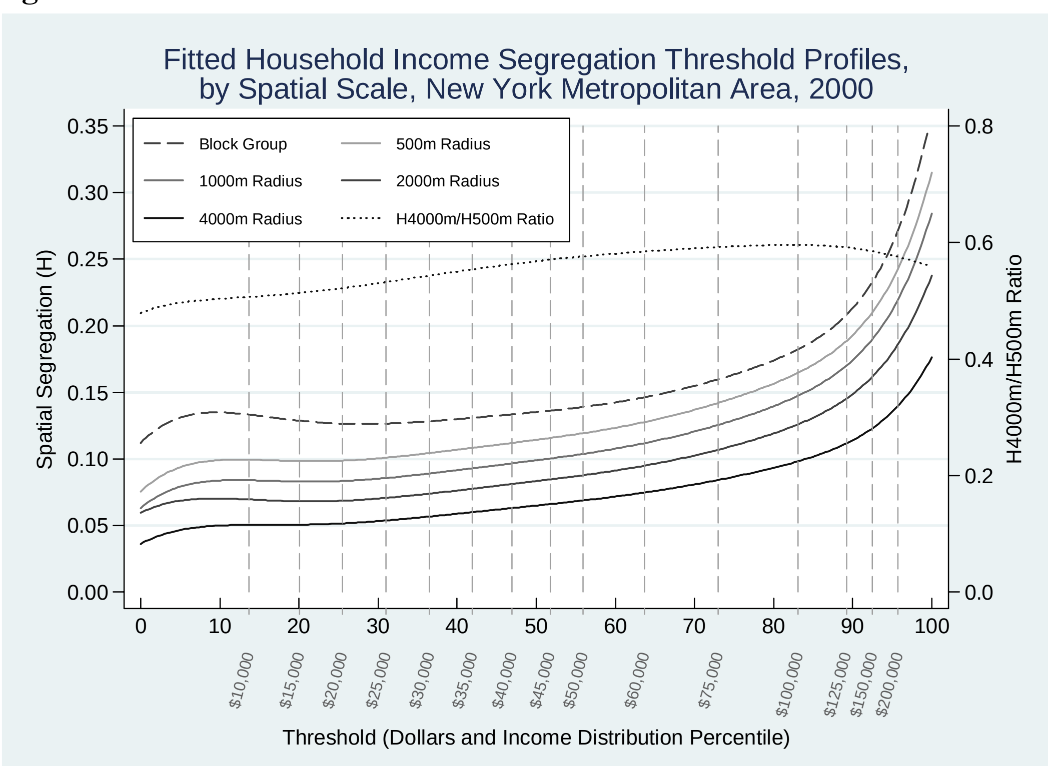

Figure 10

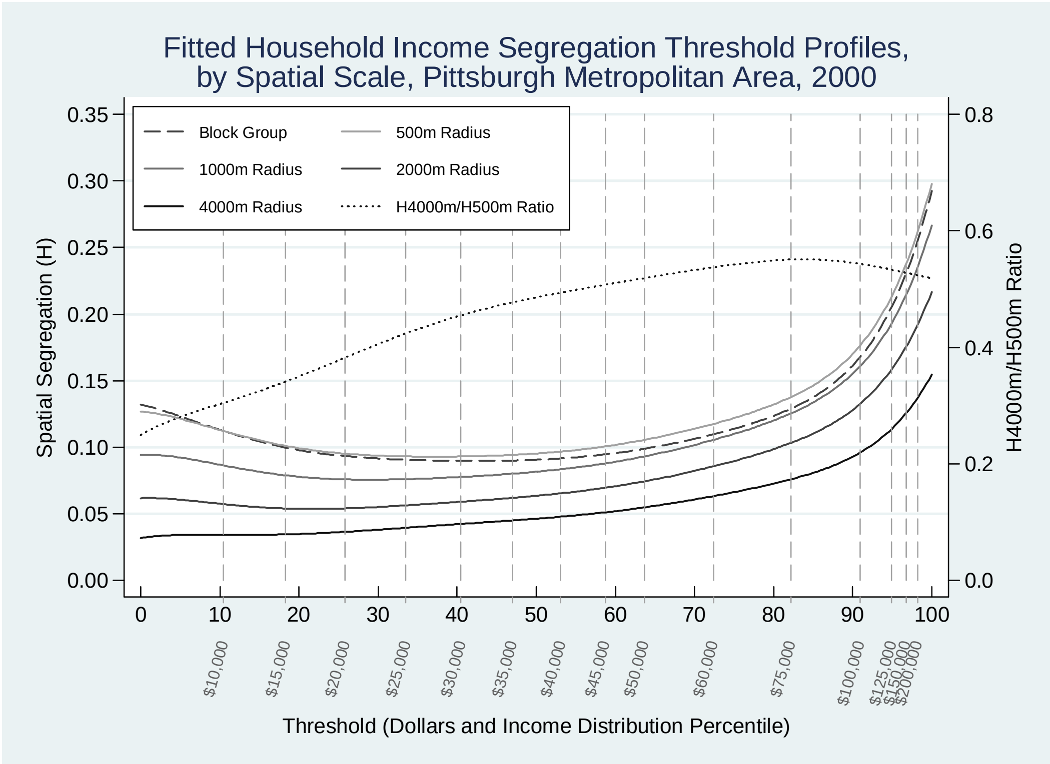

Figure 11

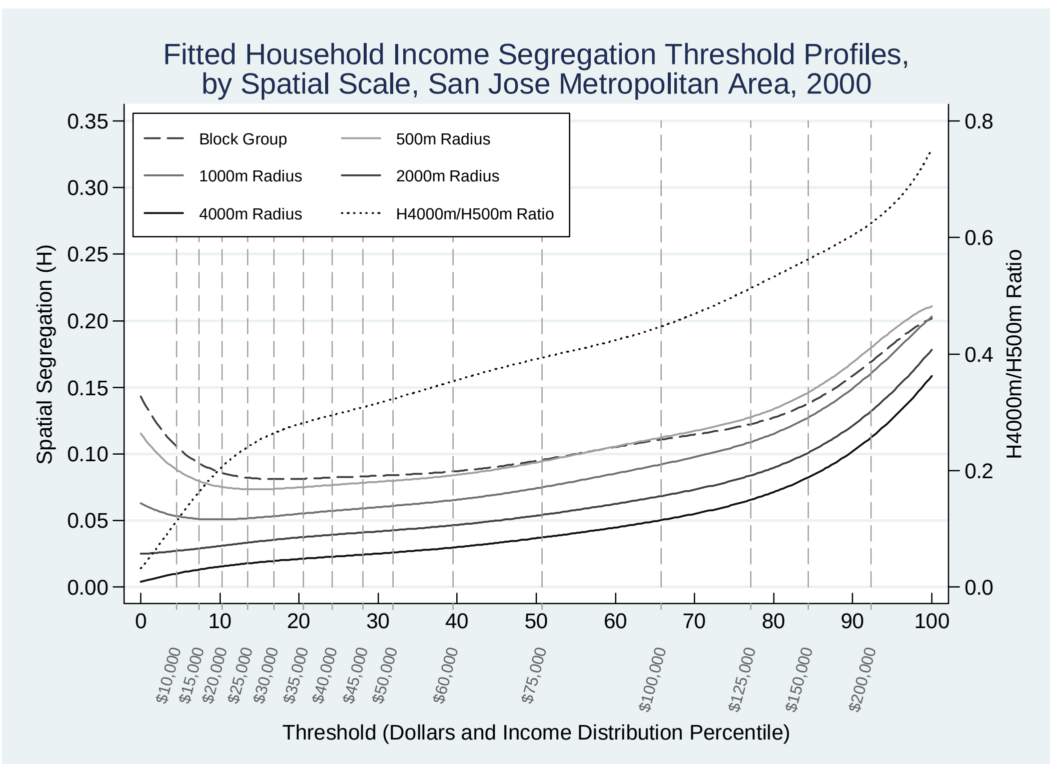

Table 1: Estimated Household Income Segregation Levels, Selected Metropolitan Areas, 2000

Panel A: Rank-Order Information Theory Index

Metropolitan Area	H_O	H_R (M=2)	H_R (M=3)	H_R (M=4)	H_R (M=5)	H_R (M=6)	H_R (M=7)	H_R (M=8)	H_R (M=9)	H_R (M=10)
Atlanta	0.1360	0.1365	0.1362	0.1362	0.1362	0.1364	0.1364	0.1364	0.1364	0.1365
Denver	0.1652	0.1670	0.1664	0.1662	0.1662	0.1664	0.1664	0.1664	0.1664	0.1666
Minneapolis	0.1385	0.1366	0.1360	0.1357	0.1357	0.1360	0.1360	0.1360	0.1360	0.1360
New York	0.1475	0.1478	0.1462	0.1472	0.1465	0.1476	0.1475	0.1466	0.1467	0.1469
Pittsburgh	0.1088	0.1062	0.1053	0.1055	0.1049	0.1055	0.1054	0.1051	0.1052	0.1052
San Jose	0.0974	0.1020	0.1020	0.1028	0.1028	0.1031	0.1030	0.1029	0.1029	0.1029

Panel B: Rank-Order Relative Diversity Index

Metropolitan Area	R_O	R_R (M=2)	R_R (M=3)	R_R (M=4)	R_R (M=5)	R_R (M=6)	R_R (M=7)	R_R (M=8)	R_R (M=9)	R_R (M=10)
Atlanta	0.1472	0.1523	0.1528	0.1529	0.1529	0.1528	0.1528	0.1528	0.1528	0.1528
Denver	0.1764	0.1846	0.1849	0.1849	0.1849	0.1848	0.1848	0.1848	0.1848	0.1849
Minneapolis	0.1465	0.1504	0.1503	0.1504	0.1504	0.1503	0.1503	0.1503	0.1503	0.1503
New York	0.1620	0.1613	0.1613	0.1611	0.1612	0.1609	0.1610	0.1608	0.1609	0.1609
Pittsburgh	0.1161	0.1159	0.1159	0.1158	0.1160	0.1159	0.1160	0.1160	0.1160	0.1160
San Jose	0.1029	0.1141	0.1143	0.1140	0.1140	0.1141	0.1141	0.1140	0.1140	0.1138

Table 2: Estimated Household Spatial Income Segregation Levels, by Spatial Scale, Selected Metropolitan Areas, 2000

Panel A: Rank-Order Information Theory Index

Metropolitan Area	**Aspatial H_R** (Block groups)	**Spatial H_R** (500m)	(1000m)	(2000m)	(4000m)	**Granularity Ratio** $H_{R(4000m)}/H_{R(500m)}$
Atlanta	0.1364	0.1457	0.1355	0.1177	0.0951	0.6528
Denver	0.1664	0.1685	0.1435	0.1143	0.0905	0.5368
Minneapolis	0.1360	0.1422	0.1229	0.0980	0.0733	0.5157
New York	0.1466	0.1236	0.1076	0.0906	0.0698	0.5646
Pittsburgh	0.1051	0.1098	0.0939	0.0726	0.0517	0.4704
San Jose	0.1029	0.1020	0.0821	0.0605	0.0431	0.4226

Panel B: Rank-Order Relative Diversity Index

Metropolitan Area	**Aspatial R_R** (Block groups)	**Spatial R_R** (500m)	(1000m)	(2000m)	(4000m)	**Granularity Ratio** $R_{R(4000m)}/R_{R(500m)}$
Atlanta	0.1528	0.1630	0.1524	0.1332	0.1082	0.6639
Denver	0.1848	0.1878	0.1614	0.1296	0.1034	0.5508
Minneapolis	0.1503	0.1577	0.1375	0.1106	0.0835	0.5296
New York	0.1608	0.1378	0.1206	0.1017	0.0782	0.5678
Pittsburgh	0.1160	0.1215	0.1046	0.0814	0.0583	0.4800
San Jose	0.1140	0.1140	0.0928	0.0688	0.0490	0.4299

Note: All values of H_R and R_R are estimated using an 8^{th}-order fitted polynomial approximation. Spatial indices are computed using a biweight kernel proximity function with bandwidths 500m, 1000m, 2000m, and 4000m.