Journal

Spatial Economic Analysis ›

This journal

- Aims and scope
- Instructions for authors
- Society information

Journal information

Print ISSN: 1742-1772 Online ISSN: 1742-1780

4 issues per year

Spatial Economic Analysis is indexed in AEA; EBSCOhost (Business Source Complete, Business Source Corporate, Current Abstracts, TOC Premier); EconLit; ERSA; NHN; OCLC (ArticleFirst, Electronic Collections Online); RePEc; Scopus and Social Sciences Citation Inde; The IET;

Article Accepted for Publication in *Spatial Economic Analysis*

This certificate confirms that the article
'Inference for the neighborhood inequality index'
by Francesco Andreoli and Eugenio Peluso
(University of Verona and LISER)
was accepted for publication in the journal on 6 July 2020.

It will be published online in 2020 and in a print issue of the journal in 2021.

Paul Elhorst
Editor-in-Chief
Spatial Economic Analysis
Print ISSN: 1742-1772
Online ISSN: 1742-1780

Tuesday, 7 July 2020

Inference for the neighborhood inequality index*

Francesco Andreoli† Eugenio Peluso

University of Verona and LISER

June 2020

Abstract

The neighborhood inequality (NI) index measures aspects of spatial inequality in the distribution of incomes within a city. The NI index is a population average of the normalized income gap between each individual's income (observed at a given location in the city) and the incomes of the neighbors located within a certain distance range. The approach overcomes the Modifiable Areal Units Problem affecting local inequality measures. This paper provides minimum bounds for the NI index standard error and shows that unbiased estimators can be identified under fairly common hypothesis in spatial statistics. Results from a Monte Carlo study support the relevance of the approximations. Rich income data are then used to infer about trends of neighborhood inequality in Chicago, IL over the last 35 years.

Keywords: individual neighborhood, variogram, ratio measures, variance approximation, income, US census and ACS.

JEL codes: C12, C46, D63, R23.

*We are grateful to Flavio Santi, three anonymous referees and the editor for their valuable comments. Errors are of course ours. This paper forms part of the research project *The Measurement of Ordinal and Multidimensional Inequalities* (grant ANR-16-CE41-0005-01) of the French National Research Agency and the projects MOBILIFE (grant RBVR17KFHX) and PREOPP (grant RBVR19FSFA) of the University of Verona Basic Research scheme, whose financial support is gratefully acknowledged. Replication code is made available on Francesco Andreoli's web-page.

†Corresponding author. Andreoli and Peluso are both at the Department of Economics, University of Verona, Via Cantarane 24, 37129 Verona, Italy, and the Luxembourg Institute of Socio-Economic Research, LISER, MSH, 11 Porte des Sciences, L-4366 Esch-sur-Alzette/Belval Campus, Luxembourg. E-mails: francesco.andreoli@liser.lu and eugenio.peluso@liser.lu.

1 Introduction

The importance of regional disparities for economic development, social and political cohesion is well-established in the literature (see Doran, Jordan and Elhorst 2018). The increasing inequality pattern registered in the US in the last few decades seems to be replicated also at a local scale within cities (Moretti 2013, Baum-Snow and Pavan 2013, Chetty and Hendren 2018). Income inequalities that arise from differences across neighborhoods, understood as areal units defined by an exogenous partition of the urban space, have received substantial attention in the literature (Massey and Eggers 1990, Jargowsky 1997, Watson 2009, Reardon and Bischoff 2011). Less evidence is available about the extent and dynamics of income inequality *within* the neighborhood (relevant contributions are Hardman and Ioannides 2004, Shorrocks and Wan 2005, Dawkins 2007, Wheeler and La Jeunesse 2008, Kim and Jargowsky 2009). The degree of inequality within the neighborhood of residence has been found to have an independent effect on important dimensions of quality of life, such as labor market attachment (Conley and Topa 2002), well-being (Ludwig et al. 2012) health (Ludwig et al. 2011, Ludwig et al. 2013, Chetty et al. 2016) and intergenerational mobility (Andreoli and Peluso 2018).

Existing approaches to the measurement of inequality within neighborhoods rely on the exogenous partition of the urban space into areal units, and are hence exposed to the Modifiable Areal Unit Problem (Openshaw 1983, Wong 2009). The MAUP can be mitigated by addressing neighborhood inequality across contiguous areas (Chakravorty 1996, Shorrocks and Wan 2005) or among neighbors in selected clusters (Hardman and Ioannides 2004). Andreoli and Peluso (2018) suggest a class of local inequality measures based on the notion

of individual neighborhood (Galster 2001), assuming spatial continuity in the underlying income distribution. A new *Neighborhood Inequality Index* (NI) is obtained following two steps of aggregation of spatial income heterogeneity recorded in individual neighborhoods. First, inequality is assessed within each individual neighborhood of a fixed (arbitrary) size. An individual neighborhood in location s gathers all income units located within the circular region of given size centered on s, and it may overlap with other individual neighborhoods of similar size but centered on different locations. Second, the resulting values of inequality within individual neighborhoods are aggregated across the relevant population. The approach guarantees robustness vis-à-vis the MAUP, insofar individual neighborhoods depend exclusively on the spatial arrangements of incomes on the map but do not rely on a specific organization of the space.

Using Census and Commuting Survey (ACS) from American MSAs, Andreoli and Peluso (2018) find that neighborhood inequality is i) high and close to citywide inequality even when the neighborhood size is less than a mile in range, ii) on the rise since 1980s and iii) displaying similar patterns across cities. The NI index estimates may be biased by measurement error and by the sampling design of spatial data, whereas bias cannot be effectively dealt with without an appropriate inference strategy.

In this paper, we derive minimum bounds for the standard error of the NI index and use these bounds to infer about robust changes in NI. We utilize some properties of the ratio estimators in Goodman and Hartley (1958) to derive bounds for the NI index variance when the data generating process is not i.i.d., accommodating for the possibility of spatial dependence. We assume that the process is continuous in nature and relies exclusively

on information about income levels and their geocoded location. We then show (Sections 2 and 3) that under fairly common assumptions in spatial statistics the estimators of the NI index standard error are identified by moments of the spatial income distribution and by the variogram, a measure of spatial dependence of the data (Matheron 1963). A simulation study (Section 4) confirms the qualities of the standard error estimator proposed here. In Section 5, we infer about changes in NI in Chicago, IL, where we find robust statistical support for rising neighborhood inequality irrespectively of the chosen size of individual neighborhoods. Section 6 concludes.

2 Measuring inequality in the neighborhood

2.1 NI index and the related literature

Consider a population of $n \geq 3$ individuals, indexed by $i = 1, ..., n$. Let y_i be the income of individual i and $\mathbf{y} = (y_1, y_2, ..., y_n)$ the sample income distribution with average $\mu > 0$. Information on incomes comes with information about their location on the city map (non-stochastic). An individual neighborhood d_i gathers n_{id} individuals living in the circular region of ray d centered on location i. If each individual occupies a separate location, there would be as many different individual neighborhoods as individuals in the city. Each individual neighborhood is characterized by an average income $\mu_{id} = \frac{\sum_{j \in d_i} y_j}{n_{id}}$, while $\Delta_i(\mathbf{y}, d) = \frac{1}{\mu_{id}} \sum_{j \in d_i} \frac{|y_i - y_j|}{n_{id}}$ ia a normalized measure of the average gap between i's income and that of her neighbors. The NI index measures the degree of inequality in the

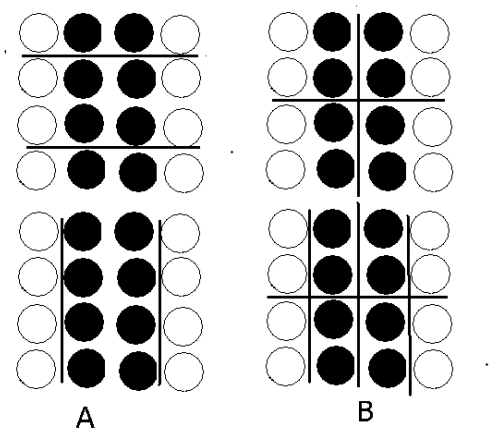

Figure 1: Scaling and zonation effects on neighborhood inequality

average individual neighborhood. It is defined as

$$NI(\mathbf{y}, d) = \frac{1}{2}\sum_{i=1}^{n}\frac{1}{n}\Delta_i(\mathbf{y}, d). \tag{1}$$

The NI index depends on d, a parameter chosen by the researcher. The plot of the value of $NI(\mathbf{y}, d)$ against d defines a *neighborhood inequality curve*. The curve is expected to be close to the origin when d goes to zero (individual neighborhoods are very small) remaining low when sample units are spatially clustered by income. When d reaches the size of the city, each individual neighborhood spans the whole city. In this case, neighborhood inequality converges to citywide inequality measured by the Gini index and the NI curve is flat.

Alternative approaches to spatial inequality within (Wheeler and La Jeunesse 2008, Shorrocks and Wan 2005) and across (Reardon and Bischoff 2011, Iceland and Hernandez 2017) areal units (such as regions, counties, census tracts, etc.) have been discussed in the literature. These approaches are subject to the Modifiable Areal Unit Problem

(Openshaw 1983, Wong 2009), that is, they are not robust to the *zonation* effect, due to the administrative partition of space (for instance, by census tracts or school districts), or to the subsequent way of *scaling* it (for instance, by block groups or by zip codes). The example in Figure 1 illustrates the implications of these effects for spatial inequality.

Each panel of Figure 1 shows the spatial distribution of rich (black dots) and poor (circles) individuals. The administrative partition of the stylized city is induced by vertical and horizontal lines. As a consequence of the zonation effect (panel A), spatial inequality measures may reverse the ranking of the two cities by just changing the design of the administrative partition (i.e. replacing the horizontal lines by the vertical ones artificially minimizes inequality within neighborhoods, inhabited by same-income individuals- and maximizes inequality between neighborhoods). As a consequence of the scaling effect (panel B), large discontinuities in spatial inequality evaluation may occur in response to minor refinements of the neighborhoods partition. In the upper diagram of Panel B, each neighborhood displays the same inequality as in the city. In the bottom diagram of Panel B, inequality within neighborhoods is eliminated by virtue of a finer partition obtained by drawing two additional vertical lines. If, instead, the finer partition were obtained by drawing horizontal lines, then the opposite situation (high inequality within the neighborhood) would have emerged.

The NI index overcomes the MAUP by treating the spatial distribution of income as continuous, and using individual neighborhoods to avoid the zonation issues. Furthermore, scaling is also controlled for by the levels of the distance threshold d, so that evaluations only depend on the actual distribution of incomes in space. These properties

guarantee that the NI index ranks as equivalent the two panels A and B in Figure 1.

Individual neighborhoods have been used in the analysis of segregation in space (Galster 2001, Clark, Anderson, Östh and Malmberg 2015), in networks (Echenique and Fryer 2007), across housing units (Hardman and Ioannides 2004) or across time (Biondi and Qeadan 2008). Chakravorty (1996) applied the notion of individual neighborhood to organizational units to develop a neighborhood disparity measure (ND), which is a normalized average of the difference between the income observed in each areal unit and the average income of neighboring parcels.

The NI index is instead related to the Gini index of inequality and develops on pairwise income comparisons within individual neighborhoods. Following Pyatt (1976), the NI index can be interpreted as the expected (relative) income gain that a randomly chosen urban resident would experience if her income was exchanged with those of her neighbors located within a ray of length d. Differently from traditional decomposable inequality measures, the NI index considers a multidimensional distribution of individual income observations alongside individual geographic locations. The two ingredients are combined through the individual neighborhoods. This allows establishing a methodological bridge with geostatistics.

2.2 Statistical properties of the NI index

Consider a spatially continuous process $\{Y_s : s = 1, \ldots, n\}$ with $s \in \mathcal{S}$ being one location on the random field, where n is the total number of locations. The process is jointly distributed as $\mathcal{F}_{\mathcal{S}}$ and may represent, for instance, the data generating process underlying

the spatial distribution of incomes across locations in a city. The joint distribution function $\mathcal{F}_{\mathcal{S}}$ combines information about the marginal income distributions in each location and the degree of spatial dependence of incomes on $\mathcal{S}$. In the analysis, locations on $\mathcal{S}$ are non-random and the process is defined conditionally on $\mathcal{S}$. Through geolocalization, it is possible to compute the distance "$||.||$" between two generic locations $s, v \in \mathcal{S}$, which we denote $||s - v|| \leq d$ or equivalently $v \in d_s$, which denotes the set of locations located within a range d from s. The cardinality of d_s is n_{d_s}. The observed spatial income distribution $\mathbf{y}$ (alongside geographic coordinates of the income observations) is a particular draw from $\mathcal{F}_{\mathcal{S}}$, where only one income realization is observed in any location s.

The NI index of the spatial process $\mathcal{F}_{\mathcal{S}}$ can be written in terms of first order moments of the random variables Y_s as follows:

$$NI(\mathcal{F}_{\mathcal{S}}, d) = \sum_{s} \sum_{v \in d_s} \frac{1}{2n\, n_{d_s}} \frac{E[|Y_s - Y_v|]}{E[Y_v]}. \tag{2}$$

The numerator in (2) depends on the extent of spatial dependence displayed by $\mathcal{F}_{\mathcal{S}}$. To show this, consider the restrictive yet widely adopted parametric assumption that Y_i follows a linear SAR model with $\mathbf{Y} = \boldsymbol{\mu} + \rho_d \mathbf{W}_d \cdot (\mathbf{Y} - \boldsymbol{\mu}) + \boldsymbol{\epsilon}$, with $\mathbf{Y} = (Y_1, \ldots, Y_n)'$, $\mathbf{W}_d$ is a $n \times n$ spatial weighting matrix with $w_{ij} = 1/n_{d_i}$ if $j \in d_i$ and $w_{ij} = 0$ otherwise, $\boldsymbol{\epsilon}$ is a column vector of i.i.d. innovations and μ the average income. The parameter ρ_d measures spatial autocorrelation at distance range d (the Moran I statistics is often used as an estimator for this correlation, see Li, Calder and Cressie 2007). Under standard assumptions (Kelejian and Prucha 2010), $\mathbf{Y} = \boldsymbol{\mu} + (\mathbf{I} - \rho_d \mathbf{W}_d)^{-1} \cdot \boldsymbol{\epsilon} = \boldsymbol{\mu} + \mathbf{S}(\rho_d) \cdot \boldsymbol{\epsilon}$, so that $Y_i = \mu + \mathbf{S}_i(\rho_d) \cdot \boldsymbol{\epsilon}$ where $\mathbf{S}_i(\rho_d)$ is a non-stochastic row vector. Under these circumstances,

$E[|Y_s - Y_v|] = E[|\mathbf{S}_s(\rho_d) \cdot \boldsymbol{\epsilon} - \mathbf{S}_v(\rho_d) \cdot \boldsymbol{\epsilon}|]$ depends on both inequality in the data ($\boldsymbol{\epsilon}$) and spatial correlation in the data ($\mathbf{S}(\rho_d)$), whereas $E[Y_v] = \mu$.

If incomes Y_s and Y_v are i.i.d. with distribution $\mathcal{F}_\mathcal{S}$ for any $s, v \in \mathcal{S}$ (i.e. $\rho_d = 0$ at any d), then $NI(\mathcal{F}_\mathcal{S}, d) = \frac{E[|Y_s - Y_v|]}{E[Y_v]}$ ($= \frac{E[|\epsilon_s - \epsilon_v|]}{\mu}$ under SAR), which coincides with a definition of the Gini index (see Muliere and Scarsini 1989). If, instead, spatial dependence is at stake (typically $\rho_d > 0$), then the NI index differs from the Gini index and $E[|Y_s - Y_v|]$ varies across locations. This quantity cannot be identified from the observation of just one data point in each location. We then introduce additional assumptions about the spatial income process that allow to derive the NI index from the first and second moments of $\mathcal{F}_\mathcal{S}$.

The first assumption is that $\mathcal{F}_\mathcal{S}$ displays *(second-order) stationarity* (see Chilès and Delfiner 2012), that is $E[Y_s] = \mu$, $Var[Y_s] = \sigma^2$ and $Cov[Y_s, Y_v] = c(||v - s||) = c(d)$ is isotropic, with $||v - s|| = d$. Under these circumstances, $Var[Y_{s+d} - Y_s] = E[(Y_{s+d} - Y_s)^2] = 2\sigma^2 - 2c(d) = 2\gamma(d)$ denotes the *variogram* of the process at distance range d (Matheron 1963). The function $2\gamma(d)$ is informative of the correlation between two random variables that are at a distance d one from the other, insofar $c(d) = \sigma^2 - \gamma(d)$. Under stationarity, $2\gamma(d) \to 0$ as d approaches 0 if the spatial process displays high positive association at small distance ranges. Conversely, $2\gamma(d) \to 2\sigma^2$ when d is sufficiently large, indicating spatial independence. We follow the convention that $s - v = d$ whenever $||s - v|| = d$ (which implies that the process occurs on a transect at neighborhood level).

The second assumption is that Y_s is Gaussian with mean μ and variance σ^2. The random variable $(Y_{s+d} - Y_s)$ is also Gaussian with variance $2\gamma(d)$, which implies $|Y_{s+d} -$

$Y_s|$ is *folded-normal* distributed (Leone, Nelson and Nottingham 1961) and its first and second moments depend exclusively on the variogram, having expectation $E[|Y_{s+d}-Y_s|] = \sqrt{2/\pi Var[Y_{s+d}-Y_s]} = 2\sqrt{\gamma(d)/\pi}$ and variance $Var[|Y_{s+d}-Y_s|] = (1-2/\pi)2\gamma(d)$.

Altogether, these assumptions allow to characterize the NI index as a function of the variogram. For given d, we consider partitioning the distance spectrum $[0,d]$ into B_d ordered intervals of size d/B_d, and derive the NI index formulation at distance intervals of fixed size. If the size was chosen equal to the minimum distance recorded between locations, then the NI index could be rewritten explicitly as a function of locations. Each interval is denoted by the index b with $b = 1,\ldots,B_d$. We further denote with d_{bi} the set of locations at interval b (and thus distant $b \cdot d/B_d$ from s_i) within the range d from location s_i. The cardinality of this set is $n_{d_{bi}} \leq n_{d_i} \leq n$. Assuming stationarity of $\mathcal{F}_\mathcal{S}$ and normality allows to write the NI index as follows:

$$\begin{aligned} NI(\mathcal{F}_\mathcal{S},d) &= \sum_i \sum_{j\in d_i} \frac{1}{2n\, n_{d_i}} \frac{E[|Y_{s_j}-Y_{s_i}|]}{\mu} \\ &= \sum_i \sum_{j\in d_i} \frac{1}{2n\, n_{d_i}} \frac{\sqrt{4\gamma(||s_j-s_i||)/\pi}}{\mu} \\ &= \sum_i \frac{1}{n} \sum_{b=1}^{B_d} \frac{n_{d_{bi}}}{n_{d_i}} \sum_{j\in d_{bi}} \frac{1}{2\, n_{d_{bi}}} \frac{\sqrt{4\gamma(s_i+b-s_i)/\pi}}{\mu} \\ &= \frac{1}{2} \sum_{b=1}^{B_d} \left(\sum_i \frac{n_{d_{bi}}}{n\, n_{d_i}} \right) \frac{\sqrt{4\gamma(b)/\pi}}{\mu}, \end{aligned} \tag{3}$$

This result (see also Andreoli and Peluso 2018), sets out the geostatistics foundations of the NI index, by showing that the index is fully characterized by the distribution of locations on the city map (non stochastic) and the degree of spatial dependence measured

by the variogram.[1] Under these assumptions, the index can be understood as an average of a standardized measure of dispersion (the second term in the summation in (3)) taken at different distance thresholds b, and weighted by the population on the random field located on the average individual neighborhood of size d, a parameter chosen by the researcher. The implications of dropping the normality assumption, widely (often implicitly) adopted in geostatistics analysis, are assessed in a Monte-Carlo experiment. Results are collected in the online appendix C.

2.3 Discussion

Few remarks are in order. First, the NI index depends both on geographical distance across locations and local income variability, averaged across all individual neighborhoods while holding distance range fixed. In presence of positive spatial correlation in incomes, small individual neighborhoods gather few likely similar income realizations, implying minimal income heterogeneity ($NI \to 0$). When the size of the individual neighborhood is large, the inequality evaluation tends to include an increasing number of income observations that are spatially unrelated. For large d, the variogram coincides with the variance of the stationary process and is constant. Equation (3) shows that the NI index stabilizes on a converging level of inequality when rising d, hence rising the weight of locations where incomes are spatially uncorrelated. This level of inequality is the Gini index of the

[1] Under second-order stationarity, the spatial autocorrelation is $\rho_d = c(d)/\sigma^2$, which implies $\gamma(d) = \sigma^2(1 - \rho_d)$. When data are i.i.d., $\gamma(d) = \sigma^2$ and NI is a local measure of inequality. Otherwise, the NI index is capable of measuring the consequences of spatial autocorrelation in the data (likely displaying $\rho_d > 0$) on local inequality estimates.

population income distribution.[2]

A second remark is about the choice of the distance parameter d. This parameter's relative magnitude is contingent to the problem under analysis. When addressing neighborhood income inequality in a urban context, it makes sense to limit the analysis to well defined statistical aggregates such as Commuting Zone or Metropolitan Statistical Areas. The geographic size of these areas is generally well defined by their administrative boundaries, density and commuting time requirements. The interpretation of the NI index, which is grounded on individual neighborhoods attributable to the residents, is meaningful when d is limited to the boundaries of the city.[3] The choice of a limit for a distance parameter may become crucial when analysing spatial inequality within inner cities, thus neglecting edge effects on the NI index and its SE bounds induced by the income distribution across the boundary areas of the city. In the online appendix D, we investigate more carefully the issue of edge effects.

Third, we stress that the NI index is a measure of inequality which is normalized by an implicit spatial weighting scheme, which we assume a uniform in (2) and (3). More precisely, each income unit observed within i's individual neighborhood is weighted $1/n_{d_i}$, whereas individual neighborhoods are weighted $1/n$. This weighting scheme allows to capture empirically aspects related to the population size, as well as the population distribution across the data field, which may cluster in dense urban areas, or sprawl in

[2]One can show that $E[|Y_s - Y_v|] = \sqrt{2Var[Y_s - Y_v]/\pi} = \sigma\sqrt{2/\pi}$ for any pair of i.i.d. normal random variables Y_s, Y_v. It follows from Muliere and Scarsini (1989) that the Gini index writes $\frac{E[|Y_s - Y_v|]}{E[Y_v]} = \sqrt{2/\pi}\frac{\sigma}{\mu}$, which is a scaled version of the coefficient of variation. In our setting, when d is large, $\gamma(d) = \sigma$. If the spatial process is i.i.d., then $NI(\mathcal{F}_{\mathcal{S}}, d)$ coincides with the Gini index at any distance threshold d.

[3]When d is large enough to include multiple urban areas, the graph of NI plotted against d may not flatten for d large. The meaningfulness of this computation rests, however, on the interest in analyzing individual neighborhoods that span over multiple urban aggregates.

suburbs. The relation between d and n in the population is described by the intensity of the population point process (the concept is formalized among others in Diggle 1985). The local density n_{d_i} is an estimator of the process intensity, which may be a source of bias in the NI index estimates. This bias vanishes when n is large[4] but may still survive for small values of d (small individual neighborhoods) if local density is too small. These concerns seem to be of secondary importance for empirical applications of the NI index in the context of urban inequality analysis. One the one hand, urban agglomerates generally display large population size and density at all distance scales. On the other hand, meaningful estimators of the empirical variogram for (3) should be based on at least 30 pairs of observations (p. 194 in Journel and Huijbregts 1989), thus providing a minimum bound for estimating the NI index at small geographic scale. Additionally, the simulation study in LeSage and Pace (2014) suggests that, at small distance ranges, the exact specification of the weighting scheme is likely irrelevant for addressing spatial variability in the data. These considerations extends to the estimation of the NI variance bounds, also based on the variogram.

3 Variance bounds for the NI index

3.1 Main result

We explore the geostatistics foundations of the NI index to derive empirically tractable bounds for the NI index variance. We denote locations on the random field with $i =$

[4]Bias can be attenuated by smoothing the observed distribution of locations using kernel estimators. As shown by Zimmerman (2008), large samples need substantially less smoothing across locations to reduce the bias, holding fixed the geographic size of the city, thus supporting the use of n_{d_i}.

$1, \ldots, n$, each having non-stochastic weight $w_i \geq 0$ with $w = \sum_i w_i$, which might reflect the inverse probability of selection from the population. The underlying process $\mathcal{F}$ is stationary with mean μ and variance σ^2.

The first implication of the assumptions we consider is that, asymptotically, the random variable $\mu_{id} = \sum_{j \in d_i} \frac{w_j}{\sum_{j \in d_i} w_j} Y_j$ is equivalent in expectation to $\tilde{\mu} = \sum_i \frac{w_i}{\sum_i w_i} Y_i$, i.e., $E[\tilde{\mu}] = \mu$. The second implication is that the spatial correlation exhibited by $\mathcal{F}$ is stationary in d and can be represented through the variogram of $\mathcal{F}$, denoted $2\gamma(d)$.

Our analysis focuses on an asymptotically equivalent version of the weighted NI index of the process distributed as $\mathcal{F}$, that is:

$$NI(\mathcal{F}, d) = \frac{1}{2\tilde{\mu}} \sum_{i=1}^{n} \sum_{j \in d_i} \frac{w_i w_j}{2\, w \sum_{j \in d_i} w_j} |Y_i - Y_j| = \frac{1}{2\tilde{\mu}} \Delta_d. \tag{4}$$

The NI index can thus be expressed as a ratio of two random variables. Asymptotic approximations for the SE of ratios of random variables have been developed in Goodman and Hartley (1958, see p. 496). Koop (1964) and Tin (1965) have demonstrated that under normality such approximations are minimum variance bounds. We use these results to obtain minimum variance bounds for the NI index in (4) as follows:

$$\begin{aligned} Var\left[NI(\mathcal{F}, d)\right] &= \frac{1}{4n\mu^2} Var[\Delta_d] + \frac{(NI(\mathcal{F}, d))^2}{n\mu^2} Var[\tilde{\mu}] - \\ &\quad \frac{NI(\mathcal{F}, d)}{n\mu^2} Cov[\Delta_d, \tilde{\mu}] + O(n^{-2}), \end{aligned} \tag{5}$$

where the SE approximation is $SE_d = \sqrt{Var\left[NI(\mathcal{F}, d)\right]}$ at any d. The approximation converges quickly when the number of locations is large, as it the case in applications

based on census micro data, and holds when income realizations are spatially correlated.[5] As suggested in Tin (1965), we use plug-in estimators for the SE.

We provide estimators for each of the three addends in (5) under appropriate assumptions. First, we assume that the process $\mathcal{F}$ is stationary with known second moments. Let the positive integer scalars m, b, b' identify intervals of the distance range d, and B is the number of such intervals. The variance of $\tilde{\mu}$, $Var[\tilde{\mu}]$, writes

$$\begin{aligned} Var[\tilde{\mu}] &= \sum_i \frac{w_i}{w} \sum_j \frac{w_j}{w} E[Y_i Y_j] - \mu^2 \\ &= \sum_i \frac{w_i}{w} \sum_{m=1}^{B} \frac{\sum_{j \in d_{mi}} w_j}{w} \sum_{j \in d_{mi}} \frac{w_j}{\sum_{j \in d_{mi}} w_j} c(||s_i - s_j||) \qquad (6) \\ &= \sum_{m=1}^{B} \left(\sum_i \frac{w_i}{w} \frac{\sum_{j \in d_{mi}} w_j}{w} (\sigma^2 - \gamma(m)) \right) \qquad (7) \\ &= \sigma^2 - \sum_{m=1}^{B} \omega(m)\gamma(m), \qquad (8) \end{aligned}$$

where (8) is obtained from (7) by renaming the weight scores so that $\sum_{m=1}^{B} \omega(m) = 1$, and by using the definition of the variogram and the fact that $s_j = s_i + m$. The score $\omega(.)$ depends on the density in one given location. In random sampling with uniform weights ($w_i = 1/n$), this factor reduces to an average of the population $n_{d_{mi}}$ residing in distance segment m from any unit i normalized by the total population, that is $\omega(m) = \frac{1}{n} \sum_i \frac{1}{n} n_{d_{mi}}$.

[5] The sample counterpart of the NI index in (4) can be interpreted as a U-statistic. As shown by Hoeffding (1948), theorem 7.5, the variance bound in (5) converges to the asymptotic unbiased estimator of the NI index variance when the income observations are i.i.d. Under this specific circumstance, asymptotic normality is also granted both with simple and with complex sampling design (Xu 2007, Davidson 2009).

The second variance component of (5), $Var[\Delta_d]$, can be written as follows:

$$\begin{aligned} Var[\Delta_d] &= \sum_{i=1}^{n}\sum_{j\in d_i}\frac{w_i w_j}{w\sum_{j\in d_i} w_j}\sum_{\ell=1}^{n}\sum_{k\in d_\ell}\frac{w_\ell w_k}{w\sum_{k\in d_\ell} w_k}E[|Y_i - Y_j||Y_\ell - Y_k|] \\ &\quad - \left(\sum_i \frac{w_i}{w}\sum_{j\in d_i}\frac{w_j}{\sum_{j\in d_i} w_j}E[|Y_j - Y_i|]\right)^2. \end{aligned}$$

The first component of $Var[\Delta_d]$ cannot be further simplified, as the absolute value operator enters the expectation term in a multiplicative way. We assume stationarity and, additionally, normality to be able to simulate the expectation, since the random vector (Y_j, Y_i, Y_k, Y_ℓ) is jointly normally distributed with expectations (μ, μ, μ, μ) and variance-covariance matrix $Cov[(Y_j, Y_i, Y_k, Y_\ell)]$, with:

$$Cov[(Y_j, Y_i, Y_k, Y_\ell)] = \begin{pmatrix} \sigma^2 & c(||s_j - s_i||) & c(||s_j - s_k||) & c(||s_j - s_\ell||) \\ & \sigma^2 & c(||s_i - s_k||) & c(||s_i - s_\ell||) \\ & & \sigma^2 & c(||s_k - s_\ell||) \\ & & & \sigma^2 \end{pmatrix}.$$

Let denote further $s_j - s_i = b \geq 0$ and $s_k - s_\ell = b' \geq 0$ for the positive integers $b \leq B_d$ and $b' \leq B_d$. We also take the (unrestrictive) convention that $s_i - s_\ell = m$ with $0 \leq m \leq B$.

We can hence express the variance-covariance matrix as a function of the variogram

$$Cov[(Y_j, Y_i, Y_k, Y_\ell)] = \begin{pmatrix} \sigma^2 & \sigma^2 - \gamma(b) & \sigma^2 - \gamma(m+b-b') & \sigma^2 - \gamma(m+b) \\ & \sigma^2 & \sigma^2 - \gamma(m-b') & \sigma^2 - \gamma(m) \\ & & \sigma^2 & \sigma^2 - \gamma(b') \\ & & & \sigma^2 \end{pmatrix}.$$

The expectation $E[|Y_i - Y_j||Y_\ell - Y_k|]$ can be simulated from a large number S (with $S = 1,000$) of independent draws $(y_{1s}, y_{2s}, y_{3s}, y_{4s})$ with $s = 1, \ldots, S$, from the random vector (Y_j, Y_i, Y_k, Y_ℓ). The simulated expectation is a function of the variogram parameters m, b, b' and d and of σ^2. It is denoted $\theta(m, b, b', d, \sigma^2)$ and estimated as follows:

$$\theta(m, b, b', d, \sigma^2) = \frac{1}{S} \sum_{s=1}^{S} |y_{2s} - y_{1s}| \cdot |y_{4s} - y_{3s}|.$$

With some algebra, and using the fact that $E[|Y_\ell - Y_i|] = 2\sqrt{\gamma(m)/\pi}$ for locations ℓ and i at distance $m \leq B$ one from each other, it is then possible to write the term $Var[\Delta_d]$ as follows:

$$\begin{aligned} Var[\Delta_d] &= \sum_{m=1}^{B} \sum_{b=1}^{B_d} \sum_{b'=1}^{B_d} \omega(m, b, b', d) \theta(m, b, b', d, \sigma^2) \\ &\quad -4 \left(\sum_{m}^{B_d} \omega(m, d) \sqrt{\gamma(m)/\pi} \right)^2 . \end{aligned} \tag{9}$$

In the formula, $\omega(m, b, b', d) = \sum_i \frac{w_i}{w} \sum_{j \in d_{bi}} \frac{w_j}{\sum_{j \in d_i} w_j} \sum_{\ell \in d_{mi}} \frac{w_\ell}{w} \sum_{k \in d_{b'\ell}} \frac{w_k}{\sum_{k \in d_\ell} w_k}$ and $\omega(m, d) = \sum_i \frac{w_i}{w} \sum_{j \in d_{mi}} \frac{w_j}{\sum_{j \in d_i} w_j}$ are calculated as before.

The third component of (5) is the covariance term. It can also be written as a function of the variogram. To show this, we maintain the convention that $s_i - s_\ell = m \geq 0$ and $s_j - s_i = b \geq 0$. This gives the following equivalence:

$$\begin{aligned} E[|Y_j - Y_i|Y_\ell] &= E[|Y_jY_\ell - Y_iY_\ell|] = E[Y_jY_\ell] - E[Y_iY_\ell] - 2E[\min\{Y_jY_\ell - Y_iY_\ell, 0\}] \\ &= c(||s_j - s_\ell||) + \mu^2 - c(||s_i - s_\ell||) - \mu^2 - 2E[\min\{Y_jY_\ell - Y_iY_\ell, 0\}] \\ &= \gamma(m) - \gamma(m+b) - 2E[\min\{Y_jY_\ell - Y_iY_\ell, 0\}]. \end{aligned} \tag{10}$$

The expectation $E[\min\{Y_jY_\ell - Y_iY_\ell, 0\}]$ is non-liner in the underlying random variables. Under the Gaussian assumption, the expectation can be simulated from a large number S (with $S = 1,000$) of independent draws (y_{1s}, y_{2s}, y_{3s}) with $s = 1, \ldots, S$, from the random vector (Y_j, Y_i, Y_ℓ), which is normally distributed with expectations (μ, μ, μ) and variance-covariance matrix $Cov[(Y_j, Y_i, Y_\ell)]$. The variance-covariance matrix writes

$$Cov[(Y_j, Y_i, Y_\ell)] = \begin{pmatrix} \sigma^2 & \sigma^2 - \gamma(b) & \sigma^2 - \gamma(m+b) \\ & \sigma^2 & \sigma^2 - \gamma(m) \\ & & \sigma^2 \end{pmatrix}$$

for given m, b and d. The resulting simulated expectation is denoted $\phi(m, b, d, \sigma^2)$ and computed as follows:

$$\phi(m, b, d, \sigma^2) = \frac{1}{S} \sum_{s=1}^{S} \min\{y_{1s}y_{3s} - y_{2s}y_{3s}, 0\}.$$

Based on this result, the covariance term in (5) becomes:

$$\begin{aligned} Cov[\Delta_d, \tilde{\mu}] &= \sum_i \frac{w_i}{w} \sum_{j \in d_i} \frac{w_j}{\sum_{j \in d_i} w_j} \sum_\ell \frac{w_\ell}{w} E[|Y_j - Y_i| Y_\ell] \\ &\quad -\mu \sum_i \frac{w_i}{w} \sum_{j \in d_i} \frac{w_j}{\sum_{j \in d_i} w_j} E[|Y_j - Y_i|] \\ &= \sum_{m=1}^{B} \sum_{b=1}^{B_d} \omega(m, b, d) \left[\gamma(m) - \gamma(m+b) - 2\phi(m, b, d, \sigma^2) \right] \\ &\quad -2\mu \sum_{m=1}^{B_d} \omega(m, d) \sqrt{\gamma(m)/\pi}. \end{aligned} \tag{11}$$

The weights in (11) coincide respectively with $\omega(m,b,d) = \sum_i \frac{w_i}{w} \sum_{\ell \in d_{mi}} \frac{w_\ell}{w} \sum_{j \in d_{bi}} \frac{w_j}{\sum_{j \in d_i} w_j}$ and $\omega(m,d) = \sum_i \frac{w_i}{w} \sum_{j \in d_{mi}} \frac{w_j}{\sum_{j \in d_i} w_j}$.

A consistent estimator for the SE, denoted $\hat{SE}_d$, is obtained by plugging into (5) the empirical counterparts of the variogram and the lag-dependent weights, using the formulas in (8), (9) and (11). (Cressie and Hawkins 1980) provide non-parametric estimators for the variogram that are robust with respect to outliers (in the empirical application, we use the robust spherical variogram model in Cressie 1985). Details are provided in the online appendix A.

3.2 Hypothesis testing

The NI index and the implied NI curves can be used to assess patterns and trends of neighborhood inequality. Various hypotheses are of interest. One concern may be about the extent at which inequality in the average individual neighborhood of size d is different from citywide inequality measured by the Gini index. The relevant null hypothesis is

$H_0^1 :\ NI(\mathbf{y}, d) = G(\mathbf{y})$ against an unrestricted alternative (reflecting the fact that neighborhood inequality can be either larger or smaller than citywide inequality). A second concern is that the empirical patterns of neighborhood inequality are related to the size of individual neighborhoods. In presence of income sorting, one expects that inequality within neighborhoods of small size to be, on average, smaller than inequality in neighborhoods of larger size. Consequently, the NI curve is expected to be increasing in the individual neighborhood size. The relevant null here is $H_0^2 :\ NI(\mathbf{y}, d') = NI(\mathbf{y}, d)$ for $d' > d$, to be tested against a restricted alternative. Rejecting both null hypotheses H_0^1 and H_0^2 gives statistical support for the existence of a neighborhood component in the urban income distribution.

It is also of interest to study the dynamics of neighborhood inequality across income distributions $\mathbf{y}_t$ and $\mathbf{y}_{t'}$. For a given size d of the individual neighborhood, the relevant null hypothesis is $H_0^3(d) :\ NI(\mathbf{y}_t, d) = NI(\mathbf{y}_{t'}, d)$ against an unrestricted one. An increase or decline in neighborhood inequality is robust (with respect to the choice of the distance parameter) when it involves a form of dominance in neighborhood inequality curves: that is, when the NI curves never intersect each others. In this case, the relevant null hypothesis is: $H_0^3 :\ H_0^3(d)\ for\ some\ d$ against a constrained alternative $\min_d\{NI(\mathbf{y}_t, d) - NI(\mathbf{y}_{t'}, d)\} > 0$, which signals that one curve lies above the other at any distance range. One particular case in which H_0^3 cannot be rejected is the situation in which $H_0^3(d)$ is true for every d. As for H_0^1 and H_0^2, the null hypotheses are expressed in the form of equalities to stress that one is compelled to conclude in favor of increasing or decreasing neighborhood inequality only if there is strong evidence against the null hypothesis.

The acceptance regions for the null hypotheses H_0^1, H_0^2 and $H_0^3(d)$ can se constructed using the confidence bounds implied by the SE approximations provided above. Confidence bounds for the NI index based on individual neighborhoods of size d take the form $\hat{NI}(\mathbf{y}, d) \pm z_\alpha SE_d$, where $\hat{NI}$ is a consistent estimator of the NI index and z_α is assumed to be the standard normal critical value for confidence level $1-\alpha$ (for instance, 95%). To test H_0^3, it is sufficient to plot the confidence bounds of $NI(\mathbf{y}_{t'}, d) - NI(\mathbf{y}_t, d)$ against d and verify that the horizontal orthant lies homogeneously in the implied rejection region. In fact, H_0^3 is rejected only if there is enough evidence against a possible crossing in NI curves, which requires to verify if the implied confidence interval bounds do not include the horizontal orthant.[6]

4 Monte Carlo study

4.1 Results

The size and power properties of the estimators adopted to test dominance in NI curves are now assessed within the framework of a Monte Carlo study. The simulation study is informative about the behaviour of the SE estimations for the NI index and the implication is has for testing null hypothesis about NI curves based on different distributions.

[6] H_0^3 cannot be rejected if there is evidence of at least some intersection between the curves. It is rejected if one curve lies everywhere above the other, implying that the former distribution displays more spatial inequality than the other irrespectively of the choice of the distance parameter. The suggested procedure of joint testing is analogous to that used in stochastic dominance analysis (Dardanoni and Forcina 1999). One can test for intersections of cdfs against the alternative of strong first order stochastic dominance by producing t-tests for the intersection of cdfs at any given income percentile (usually a grid). The null is rejected when the smallest t-test value recorded across percentiles range is larger than the corresponding 95% percentile of a standardized normal distribution (see also Bishop, Chakraborti and Thistle 1989, Andreoli 2018). We use the same logic to test for equality in NI curves taking dominance as the alternative .

Figure 2: Neighborhood inequality in Chicago, IL, 2014, versus two counterfactual distributions

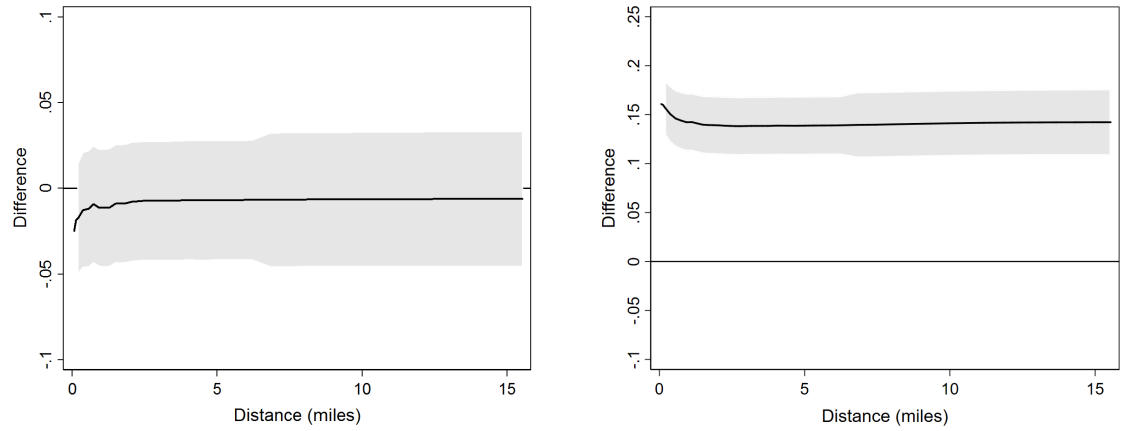

(a) $NI(\mathcal{F}_0, d) - NI(\mathcal{F}_1, d)$ (b) $NI(\mathcal{F}_0, d) - NI(\mathcal{F}_2, d)$

Note: Author analysis of US Census and ACS data. Confidence intervals are at 95% level.

The distributions are calibrated to represent the actual distribution of gross equivalent household income in Chicago IL in 2014, obtained from the Census Bureau's American Community Survey data, 2010-2014 module. We compare the actual distribution with counterfactual distributions obtained by applying suitable transformations to the actual ACS 2010-2014 data, so that these distributions can be clearly ordered in terms of NI curves dominance. Then, we use moments of these population distributions to identify moments of the income data generating processes adopted in the simulation study.

The first distribution $\mathcal{F}_0$ represents the spatial income distribution in Chicago, 2014 ($\mu_0 = \$53,456$, $\sigma_0 = \$55,310$). We further consider two counterfactual distributions $\mathcal{F}_1$ and $\mathcal{F}_2$. The distribution $\mathcal{F}_1$ is obtained by adding noise to $\mathcal{F}_0$, so that $y_1 = y_0 + \varepsilon$ for $y_1 \sim \mathcal{F}_1$, $y_0 \sim \mathcal{F}_0$ and $\varepsilon \sim N(0, 6118.44)$ ($\mu_1 = \mu_0$ and $\sigma_1 = \$55,631 > \sigma_0$). This counterfactual distribution displays similar patterns of neighborhood inequality as $\mathcal{F}_0$: H_0^3:

$NI(\mathcal{F}_0, d) = NI(\mathcal{F}_1, d)$ for at least some d, cannot be rejected, as shown in panel (a) of figure 2. The distribution $\mathcal{F}_2$ is obtained by simulating the effect of a redistributive linear income tax scheme applied to incomes distributed as $\mathcal{F}_0$, so that $y_2 = (1-t)y_0 + m$, for $y_0 \sim \mathcal{F}_0$, a flat tax rate $t = 0.3$ and basic income $m = 0.3\mu_0$ ($\mu_2 = \mu_0$, $\sigma_2 = \$38,716 < \sigma_0$). The null hypothesis H_0^3: $NI(\mathcal{F}_0, d) = NI(\mathcal{F}_1, d)$ for at least some d is clearly rejected in favor of a dominance alternative, as shown in panel (b) of figure 2.

The simulation study is based on models for the income process $\mathbf{Y}_f^n \sim (\mu_f, \sigma_f^2, \gamma_f(.))$ that is normal distributed with finite moments estimated on $\mathcal{F}_f$ with $f = 0, 1, 2$. The Monte Carlo experiment consists in randomly drawing realizations from $\mathbf{Y}_f^n$, each denoted $\mathbf{y}_{f,r}^n$ with $r = 1, \ldots, 200$, and assessing relevant nulls hypothesis at pre-determined distance threshold and for variable n. For samples of size $n = 2000,\ 5000,\ 8000$, distance cutoffs are set at approximately a third of a mile distance range increments within the first 2 miles, and then at increasing increments within the next 12 miles (at 19 miles range the NI index converges to citywide inequality). For the sample of size $n = 1000$, distance thresholds within 1 and 19 miles are set by looking at increments of three quarters of a mile exclusively. $H_0^3(d)$ is tested at each distance cutoff. The null hypothesis of the type H_0^3 is tested instead by looking at all distance cutoffs. A detailed description of the Monte Carlo study is in the online appendix.

We first investigate the *size* for the tests for various null hypothesis about NI curves. The size corresponds to the share of simulated samples that allow to reject the relevant null hypothesis when the null hypothesis is true in the population. We consider populations where H_0^3: $NI(\mathcal{F}_0, d) = NI(\mathcal{F}_1, d)$ for at least some d is true. We draw replicas $\mathbf{y}_{0,r}^n$ and

n	Distance cutoffs (miles)									# Rej.	Rej.	Weak	Strong
	0.4	0.7	1	1.4	1.7	2	3	5	12				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Panel A : Size comparisons for the true null H_0^3 **:** $NI(\mathbf{y}_{0,r}^n, d) = NI(\mathbf{y}_{1,r}^n, d),\ \forall d$													
1000	.	0.00	.	0.24	.	0.19	0.16	0.03	0.01	1.1	0.46	0.60	0.00
2000	0.00	0.00	0.32	0.22	0.23	0.19	0.10	0.08	0.00	2.4	0.62	0.53	0.00
5000	0.00	0.00	0.33	0.15	0.17	0.09	0.01	0.01	0.00	1.1	0.52	0.47	0.00
8000	0.00	0.00	0.22	0.06	0.09	0.06	0.05	0.01	0.00	0.8	0.38	0.53	0.00
Panel B : Power comparisons for the true alternative H_a^3 **:** $NI(\mathbf{y}_{0,r}^n, d) \geq NI(\mathbf{y}_{2,r}^n, d),\ \forall d$													
1000	.	0.00	.	0.29	.	0.31	0.19	0.09	0.03	1.7	0.60	0.92	0.00
2000	0.00	0.00	0.40	0.31	0.44	0.26	0.13	0.08	0.00	3.3	0.85	0.88	0.00
5000	0.00	0.00	0.55	0.28	0.49	0.24	0.12	0.09	0.03	4.3	0.81	0.98	0.00
8000	0.00	0.05	0.57	0.34	0.53	0.32	0.30	0.22	0.15	8.7	0.82	0.99	0.00

Table 1: Monte Carlo simulations of the size and power of dominance tests for NI curves that are based on the NI index SE approximations.

$\mathbf{y}_{1,r}^n$ and for each replica r we test whether $H_0^3(d)$: $NI(\mathbf{y}_{0,r}^n, d) = NI(\mathbf{y}_{1,r}^n, d)$, as well as the implied null H_0^3, are rejected by the data. Rejections are recorded and the average share of rejections over the 200 replicas is stored in Panel A of table 1. Columns (1) to (9) report the size of test for null hypothesis $H_0^3(d)$ at well defined distance cutoffs. Column (10) reports the average number of rejections of $H_0^3(d)$ across all available distance cutoffs. Column (11) reports the proportion of times that a null hypothesis $H_0^3(d)$ is rejected at least one. Columns (12) and (13) report, respectively, the share of cases where the rejection entails a weak dominance in NI curves (i.e., all cases where multiple rejections of $H_0^3(d)$ occur within the same replica r and differences in NI curves have the same sign) and the proportion of the cases in (12) where dominance is strong (i.e., H_0^3 is rejected at every distance cutoff).

Overall, the tests based on the NI index SE bounds have larger size compared to the nominal 5% level. The size of tests carried out at fixed distance cutoffs is smaller than 10% when the sample size is at least of 5000 units and it is virtually zero when $d \leq 1$ mile. Tests for $H_0^3(d)$ for $d \geq 5$ miles are below 5%. At these distance ranges, in fact,

neighborhood inequality converges to the levels of citywide inequality measured by the Gini coefficient, and the SE approximation converges asymptotically (since the spatial association of incomes becomes negligible). On average, there is less than 1 rejection of $H_0^3(d)$ across the distance cutoffs for which we test. The upper bound for the size is of 18% in the largest sample. The size of the test monotonically converges to this number as the sample size grows. A linear interpolation of size estimates in column (11) suggests that the upper bound for the size converges to its nominal value of 5% when the sample size is larger than 16,000 units.

We also investigate the *power* of the tests for various null hypothesis about NI curves. Power is measured by the share of replicas that reject the relevant null hypothesis in favor of a specific alternative when the alternative is true in the population. We use distributions so that H_0^3: $NI(\mathcal{F}_0, d) = NI(\mathcal{F}_2, d)$ for at least some d is rejected in favor of (strong) dominance in NI curves. We draw replicas $\mathbf{y}_{0,r}^n$ and $\mathbf{y}_{2,r}^n$ for replica r, and we test if $H_0^3(d)$: $NI(\mathbf{y}_{0,r}^n, d) = NI(\mathbf{y}_{2,r}^n, d)$ at each distance cutoff separately, as well as the implied null H_0^3, are rejected by the data. We find that the power of tests for $H_0^3(d)$ are relatively small for small and large distance cutoffs, while power grows above 30% for distance cutoffs between 1 and 5 miles for which we test. Tests for H_0^3 neglect the positive correlation between SE computed at different distance cutoffs, thus making rejections of the null hypothesis more likely (since part of the variability in NI curves estimates is neglected). Hence, rejections rates for H_0^3 in favor of (weak) dominance can only be interpreted as upper bounds for the power of the joint tests. These upper bounds are estimated by the product of the columns (11) and (12). The upper bound for the

power of tests for H_0^3 is of 74.8% for samples of size 2000 units and grows to 81% in the largest samples. Despite being upper bounds, these power estimates support the validity of tests for NI curves dominance based on the SE approximations even in relatively small samples. We also find that the average number of distance cutoffs where $H_0^3(d)$ is rejected at any given simulated sample grows steadily with the simulated sample size (column (10)), from 1.7 rejections when $n = 1000$ to 8.7 rejections on average when $n = 8000$, alongside larger chances that these rejections are in favor of a weak form of dominance in NI curves. Altogether, these figures confirm the relevance of the SE approximations for inferring about patterns and dynamics of neighborhood inequality.

4.2 Additional checks

In the online appendix C, we challenge the normality assumption. First, we derive SE bounds for the NI index while assuming stationarity and *log-normality* of the process as a reasonable alternative. Estimated bounds are marginally larger than those obtained under normality, implying a smaller rejection region. In the Monte Carlo experiment, we maintain the assumption that SE bounds are derived under normality, and apply these bounds to simulated data from counterfactual joint log-normal (hence non-gaussian) distributions. Compared to results in Table 1, the simulated size seems not affected by lack of normality in the underlying data, while we register improvements in simulated power estimates. This may follow from the larger acceptance region implied by the normality assumption. Overall, evidence suggests that the normality assumption leads to high-power and high-conservative tests for the null of lack of changes in neighborhood inequality.

In appendix D, we also consider the implications of edge effects arising from the choice of boundaries of the urban areas. A simulation exercise aims at assessing the implications of such effects by considering the spatial income distributions of Chicago, 2014, alongside a buffer zone on the boundaries: inequality measured in the individual neighborhoods of those living in the buffer area does not contribute to NI computation (see Griffith 1983, Xu and Dowd 2012). The sizes of the tests based on the simulated distributions (under normality) are comparable to those in table 1, but smaller for samples of 8000 units. Accounting for edge effects reduces the power of the tests for H_0^3, although the difference with table 1 is mitigated when rising sample size.

5 Inference for patterns and trends of neighborhood inequality in Chicago, IL, 1980-2014

Andreoli and Peluso (2018) provide robust evidence that neighborhood inequality is high in large American metro areas, it has been growing over the lest 35 years, and it almost converges to citywide income inequality, even when estimates are based on individual neighborhoods of small size (smaller than half a mile). Are these patterns producing reliable evidence for the population? Is the growth in neighborhood inequality statistically significant?

We use the data for the Metropolitan Statistical Area of Chicago, IL in the years 1980, 1990, 2000 and 2014 to draw inference about NI curves.[7] Chicago has experienced large

[7]See Andreoli and Peluso (2018) for details about the data. We use equivalized household market income estimates as observations, each is assumed to be located at the block group's centroid. This assumption follows the empirical empirical constraints dictated by the way data are organized into small

Figure 3: Trends in neighborhood inequality in Chicago, IL

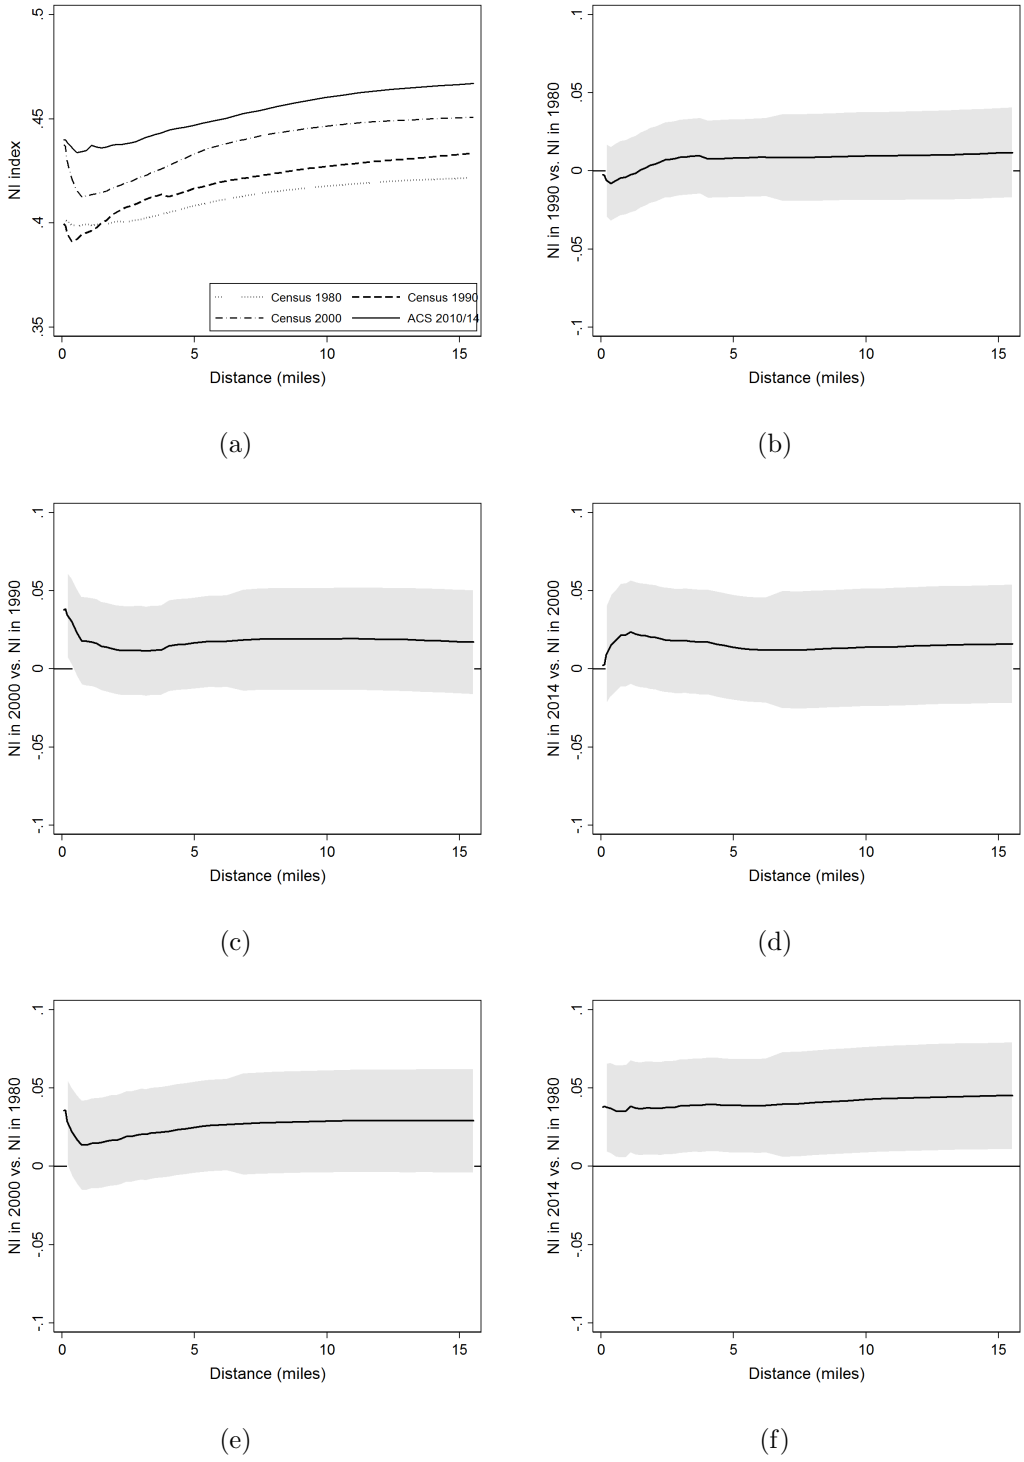

Note: Author analysis of US Census and ACS data. Confidence intervals are at 95% level.

areal units (block groups), whereas the NI index is defined for spatially continuous processes. In a simulation exercise discussed in the online appendix E we test the robustness of our estimates to such assumption.

demographic growth over the last 35 years, with the number of inhabited census blocks (each gathering approximately 1000 households) increasing from 3756 in 1980 to more than 4700 in 2014. The growth in average equivalent income (in nominal terms), ranging from \$13,794 in 1980 to \$55,710 in 2014, has been followed by an expansion of relative inequality. The Gini index for the citywide income distribution has evolved steadily, from 0.434 in 1980 to 0.461 in 1990, then to 0.473 in 2000 and finally 0.486 in 2014, reflecting both demographic and economic changes.

Neighborhood inequality in Chicago mirrors the trends observed in other large American metro areas. As shown in panel a) of figure 3, in each year the NI index is high and close to the level of the citywide Gini index even in neighborhoods of relatively small size.[8] The NI estimates are always significant at all distance ranges, with SE of the magnitude of $0.01 - 0.02$ points. As table 2 shows (panel A), the hypothesis H_0^1 is rejected with p-values always close to zero when the individual neighborhood size is smaller than 5 miles. When the individual neighborhood is of 12 miles or above, neighborhood inequality is statistically indistinguishable from the level of inequality observed in the city at conventional levels of significance in 1980, 2000 and 2014. The same table, panel B, reports the evolution of the NI index at different distance thresholds compared to the level of neighborhood inequality in individual neighborhoods of size 0.4 miles. The gap in the NI index, in italics, is positive almost everywhere and always increasing with distance. Nonetheless, these differences are not statistically significant in a distance range smaller than 5 miles. At 12 miles, H_0^2 can be rejected in every year with p-values that are slightly

[8]The nature of the Census and ACS publicly accessible data does not allow to unbiasedly estimate NI in neighborhoods smaller than 0.3 miles. Confidence intervals are only reported for larger neighborhoods.

Years	Distance d in miles					
	0.4	1	2	3	5	12
Panel A: p-values for H_0^1						
1980	0.000	0.000	0.000	0.000	0.004	0.160
1990	0.000	0.000	0.000	0.000	0.000	0.003
2000	0.000	0.000	0.000	0.000	0.001	0.070
2014	0.000	0.000	0.000	0.000	0.002	0.108
Panel B: p-values for H_0^2						
1980	.	0.493	0.454	0.396	0.239	0.067
	0	*0.000*	*0.001*	*0.003*	*0.009*	*0.020*
1990	.	0.357	0.122	0.046	0.020	0.002
	0	*0.004*	*0.014*	*0.020*	*0.025*	*0.039*
2000	.	0.311	0.410	0.461	0.239	0.060
	0	*-0.008*	*-0.004*	*0.002*	*0.012*	*0.027*
2014	.	0.467	0.465	0.390	0.269	0.071
	0	*-0.001*	*0.002*	*0.005*	*0.011*	*0.027*

Table 2: P-values for null hypothesis of the type $H_0^1 : \ NI(\mathbf{y}_t, d) = G(\mathbf{y}_t)$ and $H_0^2 : \ NI(\mathbf{y}_t, d) = NI(\mathbf{y}_t, 0.4)$, with $t = 1980,\ 1990,\ 2000,\ 2014$ and $G(\mathbf{y}_{1980}) = 0.434$, $G(\mathbf{y}_{1990}) = 0.461$, $G(\mathbf{y}_{2000}) = 0.473$, $G(\mathbf{y}_{2014}) = 0.486$. Differences in levels of the NI index are in italic.

larger than 5% (smaller in 1990). The patterns of p-values in the table confirm findings in Andreoli and Peluso (2018) that after 2000 the degree of neighborhood inequality registered in small neighborhoods has become more representative of the degree of inequality in the city.

The trends of neighborhood inequality in Chicago resemble those observed in other large American metro areas. The year-to-year changes in NI, reported in panels b), c) and d) of figure 3, are always positive at every distance range. The magnitude of these changes is, however, too small to be statistically significant according to the rejection region implied by the SE bounds. Nonetheless, the cumulated change of neighborhood inequality over the four decades turns out to be positive and significant at every distance range. As panel f) shows, the acceptance region for H_0^3 is always positive and never

includes the horizontal axis, implying that we have strong statistical evidence that the NI curve of Chicago for 2014 lies always above that of 1980 and the gap between these two curves is different from (in fact, larger than) zero.

6 Concluding remarks

This article provides variance bounds for the neighborhood inequality index. These bounds are identified from the knowledge of the variogram function which, under assumptions on the income generating process that are common in spatial statistics literature, fully characterizes the spatial income distribution.

An application to rich income data from the American Census and the Community Survey motivates the interest in using SE approximations for the NI index when assessing patterns and trends of neighborhood inequality across American cities. Focussing on the city of Chicago, IL, we find robust statistical evidence that neighborhood inequality is large even for individual neighborhoods of small size, but it is statistically different from citywide inequality (as measured by the Gini index). The cumulated growth of neighborhood inequality over the period 1980-2014 is substantial and significant at standard confidence levels, reflecting a general trend in largest American cities documented in Andreoli and Peluso (2018). Patterns are robust to granularity of the spatial distribution as well as to edge effects. The Monte Carlo study shows that the tests for NI curves dominance based on the SE approximations have higher size than the nominal values, although the (upper bound) size estimates quickly converges when the sample size grows. We expect that a sample of 16,000 units, smaller than the sample used to obtain estimates

on the 5-years ACS module, is sufficient to guarantee that the size of the tests we consider converge to their nominal values. The power of these tests is relatively small for null hypotheses defined at given distance cutoffs (but larger than 30% in simulated samples of at least 8000 units), but power grows significantly to more than 80% when considering tests for NI curves (weak) dominance (although these are only upper bounds). Investigations about the appropriate testing procedure when placing dominance/non-dominance of NI curves under the null are also left for future research.

An interesting extension is to craft individual neighborhoods based on a *nearest-neighbor* logic, that is by holding n_{d_i} as fixed and d variable. Cressie (1991) provides estimators for the variogram based on nearest-neighbor logic. While the nearest-neighborhood weighting scheme bears little practical implications for the NI calculation at small geographic scale (LeSage and Pace 2014), it fails basic replication invariance properties. Consider, for instance, the possibility of "replicating" the population so that each income unit's replica has the same income and location as the original. This operation doubles the population size and density of a city, without affecting the patterns of spatial income inequality. The NI index we study in this paper, which is normalized by population density, is not affected by this transformation. Conversely, a measure of neighborhood inequality based on the nearest-neighborhood logic would artificially dilute inequality evaluations over a larger number of neighbors. These arguments provide further support for the use of a distance-based criteria, such as the individual neighborhood, in spatial inequality measurement.

References

Andreoli, F. (2018). Robust inference for inverse stochastic dominance, *Journal of Business & Economic Statistics* **36**(1): 146–159.

Andreoli, F. and Peluso, E. (2018). So close yet so unequal: Neighborhood inequality in American cities, *ECINEQ Working paper 2018-477* .

Baum-Snow, N. and Pavan, R. (2013). Inequality and city size, *The Review of Economics and Statistics* **95**(5): 1535–1548.

Biondi, F. and Qeadan, F. (2008). Inequality in paleorecords, *Ecology* **89**(4): 1056–1067.

Bishop, J. A., Chakraborti, S. and Thistle, P. D. (1989). Asymptotically distribution-free statistical inference for generalized Lorenz curves, *The Review of Economics and Statistics* **71**(4): pp. 725–727.

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

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects*, *The Quarterly Journal of Economics* **133**(3): 1107–1162.

Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A. and Cutler, D. (2016). The association between income and life expectancy in the United States, 2001-2014., *The Journal of the American Medical Association* **315**(14): 1750–1766.

Chilès, J.-P. and Delfiner, P. (2012). *Geostatistics: Modeling Spatial Uncertainty*, John Wiley & Sons, New York.

Clark, W. A. V., Anderson, E., Östh, J. and Malmberg, B. (2015). A multiscalar analysis of neighborhood composition in Los Angeles, 2000-2010: A location-based approach to segregation and diversity, *Annals of the Association of American Geographers* **105**(6): 1260–1284.

Conley, T. G. and Topa, G. (2002). Socio-economic distance and spatial patterns in unemployment, *Journal of Applied Econometrics* **17**(4): 303–327.

Cressie, N. (1985). Fitting variogram models by weighted least squares, *Journal of the International Association for Mathematical Geology* **17**(5): 563–586.

Cressie, N. A. C. (1991). *Statistics for Spatal Data*, John Wiley & Sons, New York.

Cressie, N. and Hawkins, D. M. (1980). Robust estimation of the variogram: I, *Journal of the International Association for Mathematical Geology* **12**(2): 115–125.

Dardanoni, V. and Forcina, A. (1999). Inference for Lorenz curve orderings, *Econometrics Journal* **2**: 49–75.

Davidson, R. (2009). Reliable inference for the Gini index, *Journal of Econometrics* **150**(1): 30 – 40.

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

Diggle, P. (1985). A kernel method for smoothing point process data, *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **34**(2): 138–147.

Doran, J., Jordan, D. and Elhorst, P. (2018). Virtual special issue on regional inequality, *Spatial Economic Analysis* **13**(4): 383–386.

Echenique, F. and Fryer, R. G. (2007). A measure of segregation based on social interactions, *Quarterly Journal of Economics* **122**(2): 441–485.

Galster, G. (2001). On the nature of neighbourhood, *Urban Studies* **38**(12): 2111–2124.

Goodman, L. A. and Hartley, H. O. (1958). The precision of unbiased ratio-type estimators, *Journal of the American Statistical Association* **53**(282): 491–508.

Griffith, D. A. (1983). The boundary value problem in spatial statistical analysis*, *Journal of Regional Science* **23**(3): 377–387.

Hardman, A. and Ioannides, Y. (2004). Neighbors' incom distribution: Economic segregation and mixing in US urban neighborhoods, *Journal of Housing Economics* **13**(4): 368–382.

Hoeffding, W. (1948). A class of statistics with asymptotically normal distribution, *The Annals of Mathematical Statistics* **19**(3): 293–325.

Iceland, J. and Hernandez, E. (2017). Understanding trends in concentrated poverty: 1980-2014, *Social Science Research* **62**: 75 – 95.

Jargowsky, P. A. (1997). *Poverty and Place: Ghettos, Barrios, and the American City*, New York: Russell Sage Foundation.

Journel, A. G. and Huijbregts, C. J. (1989). *Mining geostatistics*, Academic Press, London.

Kelejian, H. H. and Prucha, I. R. (2010). Specification and estimation of spatial autoregressive models with autoregressive and heteroskedastic disturbances, *Journal of Econometrics* **157**(1): 53 – 67. Nonlinear and Nonparametric Methods in Econometrics.

Kim, J. and Jargowsky, P. A. (2009). *The Gini-coefficient and segregation on a continuous variable*, Vol. Occupational and Residential Segregation of *Research on Economic Inequality*, Emerald Group Publishing Limited, pp. 57 – 70.

Koop, J. C. (1964). On an identity for the variances of a ratio of two random variables, *Journal of the Royal Statistical Society. Series B (Methodological)* **26**(3): 484–486.

Leone, F. C., Nelson, L. S. and Nottingham, R. B. (1961). The folded normal distribution, *Technometrics* **3**(4): 543–550.

LeSage, J. P. and Pace, R. K. (2014). The biggest myth in spatial econometrics, *Econometrics* **2**(4): 217–249.

Li, H., Calder, C. A. and Cressie, N. (2007). Beyond moran's i: Testing for spatial dependence based on the spatial autoregressive model, *Geographical Analysis* **39**(4): 357–375.

Ludwig, J., Duncan, G. J., Gennetian, L. A., Katz, L. F., Kessler, R. C., Kling, J. R. and Sanbonmatsu, L. (2012). Neighborhood effects on the long-term well-being of low-income adults, *Science* **337**(6101): 1505–1510.

Ludwig, J., Duncan, G. J., Gennetian, L. A., Katz, L. F., Kessler, R. C., Kling, J. R. and Sanbonmatsu, L. (2013). Long-term neighborhood effects on low-income families: Evidence from Moving to Opportunity, *American Economic Review* **103**(3): 226–31.

Ludwig, J., Sanbonmatsu, L., Gennetian, L., Adam, E., Duncan, G. J., Katz, L. F., Kessler, R. C., Kling, J. R., Lindau, S. T., Whitaker, R. C. and McDade, T. W. (2011). Neighborhoods, obesity, and diabetes - A randomized social experiment,

New England Journal of Medicine **365**(16): 1509–1519. PMID: 22010917.

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

Matheron, G. (1963). Principles of geostatistics, *Economic Geology* **58**(8): 1246–1266.

Moretti, E. (2013). Real wage inequality, *American Economic Journal: Applied Economics* **5**(1): 65–103.

Muliere, P. and Scarsini, M. (1989). A note on stochastic dominance and inequality measures, *Journal of Economic Theory* **49**(2): 314 – 323.

Openshaw, S. (1983). *The modifiable areal unit problem*, Norwick: Geo Books.

Pyatt, G. (1976). On the interpretation and disaggregation of Gini coefficients, *The Economic Journal* **86**(342): 243–255.

Reardon, S. F. and Bischoff, K. (2011). Income inequality and income segregation, *American Journal of Sociology* **116**(4): 1092–1153.

Shorrocks, A. and Wan, G. (2005). Spatial decomposition of inequality, *Journal of Economic Geography* **5**(1): 59–81.

Tin, M. (1965). Comparison of some ratio estimators, *Journal of the American Statistical Association* **60**(309): 294–307.

Watson, T. (2009). Inequality and the measurement of residential segregation by income in American neighborhoods, *Review of Income and Wealth* **55**(3): 820–844.

Wheeler, C. H. and La Jeunesse, E. A. (2008). Trends in neighborhood income inequality in the U.S.: 1980–2000, *Journal of Regional Science* **48**(5): 879–891.

Wong, D. (2009). The modifiable areal unit problem (MAUP), *The SAGE handbook of spatial analysis* pp. 105–124.

Xu, C. and Dowd, P. A. (2012). *The Edge Effect in Geostatistical Simulations*, Springer Netherlands, Dordrecht, pp. 115–127.

Xu, K. (2007). U-statistics and their asymptotic results for some inequality and poverty measures, *Econometric Reviews* **26**(5): 567–577.

Zimmerman, D. L. (2008). Estimating the intensity of a spatial point process from locations coarsened by incomplete geocoding, *Biometrics* **64**(1): 262–270.