

Inequality at a Location: A Local Inequality Index Based on Mean Log Deviation

Dongyuan Mu

touen.mu@gmail.com

Kindai University

Research Article

Keywords:

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9134754/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Inequality at a Location: A Local Inequality Index Based on Mean Log Deviation

Abstract

Inequality is usually measured for a population, yet it is experienced from a location. This paper treats the inequality at a location as an object of measurement. We propose two properties that a measure of it should satisfy, Locality and Spatiality, and show that the existing spatial and local measures we examine satisfy at most one. We then construct the Local Inequality Index (LII), which weights each unit in the mean log deviation by a distance-decay kernel centred on the focal location, so that position is read by the measure itself rather than used to select a sample. For any decay parameter $\alpha > 0$, the LII satisfies both properties while remaining an inequality measure, nests the global MLD, and, under a homogeneous kernel, responds to the decay parameter through a covariance between local income deviations and distance. An application to Chinese prefecture-level cities from 2000 to 2020 reveals isolated high-inequality agglomerations and regions where local disparities are quietly worsening, patterns that the conventional MLD leaves entirely masked.

1 Introduction

Inequality has conventionally been regarded as a property of a sample, a statistic on the same footing as its mean or variance. Yet inequality is experienced from somewhere. Theories of relative deprivation hold that individuals assess their position against a reference group rather than against the population at large (Yitzhaki 1979), and empirical work finds that this reference group is drawn largely from their surroundings: the incomes people observe nearby shape their well-being, their perceived place in the distribution, and their views on inequality and redistribution (Luttmer 2005; Cruces, Perez-Truglia, and Tetaz 2013; Newman, Johnston, and Lown 2015; Domènech-Arúmi 2025). Local inequality also matters for other outcomes: it has been linked to crime (Kelly 2000), the discontent of places left behind has become a force in national politics (Rodríguez-Pose 2018), and place-based policy is designed around where disadvantage sits rather than how much of it there is in total (Austin, Glaeser, and Summers 2018). What matters for these phenomena is the inequality at a location, the dispersion of incomes that a place faces in its surroundings, rather than the inequality of a population.

Measuring the inequality at a location has so far meant computing a conventional index within a local area. Much empirical work computes the Gini coefficient or a similar index within small administrative units such as counties (De Nicolò and Ferrante 2025; Kelly 2000), and more recent local measures apply such indices within a fixed-radius buffer around each location (Andreoli and Peluso 2021; Domènech-Arúmi 2025). In both cases a population measure is reproduced on a spatial subsample: position is only used to decide which units enter, and plays no further role. The value therefore depends on where the boundary is drawn, on both its

scale and its delineation, which is known as the modifiable areal unit problem (MAUP). What is missing is a measure built for the object itself, one that is taken at a location and reads the position of every unit.

This paper makes three contributions. First, it turns the inequality at a location from an intuition into a criterion. We propose two properties that a measure of it should satisfy: Locality, that the measure is taken at a location and can differ across locations, and Spatiality, that it keeps reading where units stand, so that moving a unit changes the measure somewhere, however small the neighbourhood considered. Together they form a standard against which any spatial or local measure of inequality can be checked, so that whether a measure captures the inequality at a location, or merely applies a population measure to a local sample, becomes a matter of verification rather than of labelling.

Second, we construct a measure that meets this standard, the Local Inequality Index (LII). Every unit still enters the index, but a kernel function adjusts the weights and the local reference mean. As a result, the reference mean becomes the income environment visible from there, and a disparity in close surroundings weighs more than an equally large one far away. Position is thus read by the measure itself rather than used to select a sample, and for $\alpha > 0$ the LII satisfies both Spatiality and Locality while remaining an inequality measure. A practical consequence is granularity. A conventional index computed from city-level data describes only a group of cities, such as a province; the LII returns a value for every single city from the same data, describing the income environment around each city at the resolution at which the data are observed.

Third, we recommend a homogeneous kernel for the LII. In this case, the decay parameter α governs how widely the reference environment of each location is

drawn. The LII depends on geography only through relative distances; starting from the global MLD at $\alpha = 0$, the values at different locations diverge as α rises, each moving according to a covariance between local income deviations and log distance, and as $\alpha \rightarrow \infty$ each contracts to its immediate neighbourhood. We apply the measure to Chinese prefecture-level cities from 2000 to 2020. Mapped across space, the LII shows where local inequality is high, which neither a national aggregate nor a regional decomposition can reveal.

The rest of the paper is organized as follows. Section 2 defines the inequality at a location, derives its two properties and places existing measures against them. Section 3 constructs the LII, and Section 4 establishes its properties. Section 5 presents the empirical application, and Section 6 concludes.

2 Inequality at a Location

An inequality measure is concerned with how a given aggregate income is divided among the population: what matters is the distribution of shares, not their scale. The population enters the measure as a list, with no notion of whether one is standing nearer to or farther from another, and nothing in this list corresponds to location. Two configurations with the same shares, one where rich and poor sit side by side and one where they live at opposite ends of the country, return the same value.

This paper addresses this limitation by turning from the inequality of a population to the inequality at a location. The inequality at a location is the inequality one would feel standing there and looking around: not the gap between it and any one specific neighbour, but the dispersion of the environment one can see from

that location. This is not a list-like environment, one and the same for everyone; it is a location-specific income distribution, centred on that location and looking outward. Every other location is seen, but not equally: each enters with a weight that shrinks as distance grows. For instance, a poor unit may know that a wealthy unit lies a thousand miles away, but that disparity, being so remote, weighs little on the unit's own inequality, while an ordinary disparity nearby weighs far more. We use this as a modelling device: the measure records the dispersion of the reference income distribution seen from a location, not a behaviourally validated index of perceived inequality or relative deprivation.

We draw two properties from this description, and together they separate a measure of inequality at a location from a measure of inequality in a population. First, the inequality at a location is predicated of a place rather than of a population. The environment described previously is centred on the location one stands at, so that every location looks out on a different environment and receives a value of its own.

Property 1 (Locality). The measure is defined at each location and depends on it: in some configurations with incomes not all equal, it takes different values at different locations.

The second property comes from the statement that every other location is seen as a location rather than as an income alone: what the measure reads is the spatial configuration of the units, where each of them stands relative to every other. Spatial information may be brought in by more than one device, and nothing in the description requires a particular one. What it does require is that position never stops being read. If a unit could be shifted a little without any effect, then over

that range its exact position would carry no information, and what the measure records of it is not where it stands but only which side of something it falls on.

Property 2 (Spatiality). Incomes are attached to locations, and the measure responds to their positions: for any configuration of at least three locations in general position with incomes not all equal, and any unit, however small a neighbourhood of its position is taken, some position within it is such that moving the unit there changes the measure at some location.¹

A configuration is in general position if every unit has a unique nearest neighbour and no unit is the nearest neighbour of every other unit; the second condition requires at least three locations. The two properties are of different kinds. Locality states what sort of object a measure of inequality at a location is, and Spatiality states how such an object must respond to the arrangement of the units. Neither implies the other.

Table 1 places the main spatial and local inequality measures in the literature against these two properties.

Conventional measures and regional decomposition (Theil 1967; Akita 2003; Shorrocks and Wan 2005) are the trivial cases of the comparison. Both return a single statistic for the population, so neither is local. Regional decomposition does use spatial information, but only to divide the sample into groups. The groups are then treated as a list, not as units with locations. Moving a unit within its group changes nothing, so neither is spatial.

More recent work decomposes the location-related part of overall inequality (Rey and Smith 2013; Márquez, Lasarte-Navamuel, and Lufin 2018). The output

¹The index proposed in this paper does not satisfy this property in the two-location case.

Table 1: Existing measures relative to the two properties

Measure	Representative studies	Spatiality (Property 2)	Locality (Property 1)
Conventional inequality measures	(Theil 1967)	×	×
Regional decomposition	(Shorrocks and Wan 2005; Akita 2003)	×	×
Spatial decomposition of inequality measures	(Rey and Smith 2013; Márquez, Lasarte-Navamuel, and Lufin 2018)	$\triangle$	×
Neighbourhood inequality (NI)	(Andreoli and Peluso 2021)	×	×
Local neighbourhood Gini (LNG)	(Domènech-Arumí 2025)	×	✓
Local inequality index (this paper)		✓	✓

Notes: ✓ the property holds, × it does not, $\triangle$ the verdict depends on the spatialisation device the construction leaves open. For a global measure, “the measure at some location” in Property 2 refers to its single value. The entries record whether each measure has the properties of the object defined above, not an assessment of the measure in relation to the question it was designed to answer.

is a single number for the whole, so these measures are not local. Both can be run with a binary adjacency indicator or with weights that vary continuously with distance,² and only the latter satisfies Spatiality. This stream is spatial only under that choice of weights, and it is never local.

Some local measures compute inequality among the observations falling within a spatial buffer around a location (Domènech-Arú 2025). These are local measures but not spatial ones. The procedure amounts to estimation after resampling: once the units inside the buffer have been selected, moving any of them within the buffer leaves the value unchanged.

The neighbourhood inequality index of Andreoli and Peluso (2021) is built bottom-up. For each unit it first computes a local, standardized measure of dispersion from a buffer around it, the gap between that unit’s income and its neighbours’, then aggregates these local values into a single number. It is this aggregation step that removes locality. NI fails Spatiality because weighting within a buffer is uniform: moving a unit within the buffer leaves the value unchanged.

3 Construction of Local Inequality Index

3.1 Setup

Consider a finite collection of n spatial units indexed by i , j and t , each ranging over $\{1, \dots, n\}$. Throughout, j denotes the focal unit at which a local index is evaluated, while i and t are running indices over all units. Unit i holds $N_i > 0$

²In Rey and Smith (2013), an inverse distance matrix may replace the binary weights once row standardized.

income recipients³ whose mean income is $\bar{X}_i > 0$, so that $N = \sum_i N_i$ is total population and $\bar{X} = \sum_i N_i \bar{X}_i / N$ is the overall mean income. The distance between units i and j is denoted d_{ij} ; it is assumed symmetric, $d_{ij} = d_{ji}$, and strictly positive for $i \neq j$, the self-distance d_{jj} being treated in Section 3.2. When units are areas, $\bar{X}_i$ assigns every resident of i the same income, so the index measures dispersion across areas; within-area dispersion is not captured.

Weights are assigned by a kernel $w(\cdot)$, which enters only through distance: unit i carries weight $w(d_{ij})$ in the neighbourhood of j . We require of w only that it be non-negative and that $\sum_t w(d_{jt})N_t > 0$ for every j , so that the shares defined below exist. Among the GE family, we adopt the mean log deviation (MLD), whose global form is

$$MLD = - \sum_i p_i \ln \frac{s_i}{p_i} \quad (1)$$

where $p_i = N_i/N$ is the population share of unit i and $s_i = \bar{X}_i N_i / (\bar{X} N)$ is its income share. A weighted Gini can also be built from a location-specific reference distribution, so the choice of MLD is not forced by Properties 1 and 2. We choose MLD because it equals the divergence $\sum_i p_i \ln(p_i/s_i)$ between population shares and income shares, is monotonically related to the Atkinson index with $\varepsilon = 1$, and is additively decomposable; these interpretations carry over to the local version below. Writing $S_j(\cdot)$ for the resulting spatial weighting operator, the spatially weighted shares are

$$S_j(p_i) = \frac{w(d_{ij})N_i}{\sum_t w(d_{jt})N_t}, \quad S_j(s_i) = \frac{w(d_{ij})N_i \bar{X}_i}{\sum_t w(d_{jt})N_t \bar{X}_t} \quad (2)$$

which are non-negative and sum to one, since summing the numerator over i

³A unit is a single individual when $N_i = 1$.

reproduces the denominator. Replacing p_i and s_i in Equation (1) by these shares gives the LII at j :

$$LII_j := - \sum_i S_j(p_i) \ln \frac{S_j(s_i)}{S_j(p_i)} = - \sum_i S_j(p_i) \ln \frac{\bar{X}_i}{\bar{X}_j^w} = \ln \bar{X}_j^w - \sum_i S_j(p_i) \ln \bar{X}_i, \quad (3)$$

where $\bar{X}_j^w = \sum_t w(d_{jt}) N_t \bar{X}_t / \sum_t w(d_{jt}) N_t$ is the spatially weighted mean income centred on j .

The intuition behind the measure is straightforward: LII_j compares each unit i 's income $\bar{X}_i$ against the local reference mean $\bar{X}_j^w$, so the inequality that j faces is defined relative to its own income environment rather than a global average. Moreover, the kernel $w(d_{ij})$ adjusts the weight in the summation, so units closer to j not only shape the local reference mean $\bar{X}_j^w$ more heavily but also contribute more to the aggregation of log-deviations. The income disparities closer to j therefore matter more than remote large disparities.

The location-specificity of LII_j follows directly from its construction, and it is what gives LII the two properties of Section 2. Regarding Locality, because the kernel $w(d_{ij})$ is indexed by j , both the spatial weights and the weighted mean $\bar{X}_j^w$ shift whenever j changes, so that, for a kernel that is not constant, the values in general differ across units. Regarding Spatiality, a kernel that varies with distance rather than treating every unit within some cutoff alike ensures that some displacement of a unit changes the weight and the local reference mean, and hence LII_j at some location. Proposition 5 makes this precise for the homogeneous kernel. Taking, in addition, a distance-decay kernel gives the intuition that closer disparities weigh more heavily than remote ones of the same size, consistent with the literature on perceived and reference-dependent inequality (Luttmer 2005).

Two components of the LII remain to be specified: (1) the distance of a unit to itself d_{jj} , and (2) the kernel function $w(\cdot)$. We discuss these two components separately in the following subsections.

3.2 Self-distance

A spatial weight matrix, which assigns a smaller weight to more distant units, conventionally sets its diagonal elements $w(d_{jj})$ to zero, which excludes the focal unit from the calculation. This convention is appropriate for measures of spatial association such as Local Moran's I (Anselin 1995) and spatial lag. Some settings, however, require the focal unit itself to be taken into account, and this calls for a strictly positive self-distance. The same requirement arises in the accessibility literature, where the contribution a zone makes to its own accessibility is known as its self-potential and cannot be evaluated without assigning that zone an internal distance (Condeço-Melhorado et al. 2016). The LII measures the inequality of the income distribution environment that unit j faces in its surroundings, and j 's own income is a constituent part of that environment. Excluding $i = j$ would discard this information, and would additionally break the exact nesting of the global MLD established in Proposition 4(ii)(a).

Retaining $i = j$ requires a positive self-distance for the homogeneous kernel adopted below, since $d_{jj} = 0$ leaves $d_{jj}^{-\alpha}$ undefined for $\alpha > 0$. We define it in proportion to the distance to the nearest other unit, following the convention used in the gravity literature to measure distances internal to a country, where values of

k between one quarter and one half have been adopted (Head and Mayer 2010):

$$d_{jj} = k \cdot \min_{i \neq j} d_{ij}, \quad 0 < k \leq 1. \quad (4)$$

The parameter k fixes how much closer to j its own income is treated as being than that of its nearest neighbour; the boundary value $k = 1$ is admitted because it is the critical convention in Proposition 4(ii)(b). This assigns j the largest single kernel weight in its own neighbourhood, weakly so at the boundary value $k = 1$, where j ties with its nearest neighbour, consistent with the premise that the closest disparities are felt most acutely, while keeping the weight finite. This setup leads to a consequence worth noting for geographically isolated units: their weights are concentrated on themselves and on their nearest neighbours, since their remaining neighbours are distant enough for the kernel to discount them heavily. This is a faithful representation of the construction rather than a distortion, as the income environment such a unit faces is indeed composed of itself and a small number of neighbours.

With the homogeneous kernel, the weight of j relative to its nearest neighbour i is $S_j(p_j)/S_j(p_i) = (N_j/N_i) k^{-\alpha}$, so k and α jointly determine how much of j 's environment is j itself. The limit $LII_j \rightarrow 0$ for $k < 1$ as $\alpha \rightarrow \infty$ is therefore a consequence of the self-distance rule, not evidence that inequality vanishes. The rule also depends on how units are drawn: splitting a unit into two nearby parts with the same income changes d_{jj} and hence the index, so the construction reduces cut-off sensitivity but does not remove the MAUP arising from the choice of spatial units. Neither α nor k is pinned down by the construction: α expresses an assumption about how quickly disparities lose salience with distance, and k how much weight

a unit places on itself relative to its surroundings. Both are left to the researcher and should suit the setting at hand.

3.3 Kernel function

The kernel $w(\cdot)$ is where the substantive assumption about spatial perception enters the construction. It specifies how rapidly a disparity loses salience with distance, and this is a claim about the environment being measured rather than a technical device to be calibrated against the data.⁴

In this paper we recommend homogeneous kernels of the form

$$w(d) = d^{-\alpha}, \quad \alpha \geq 0, \quad (5)$$

for three reasons. First, this specification carries through, in a concrete functional form, the intuition that closer disparities should weigh more heavily on j 's LII than remote ones, as urged by Luttmer (2005). Second, it ensures every unit contributes with a strictly positive weight, avoiding the information loss and boundary sensitivity of fixed-radius constructions (Andreoli and Peluso 2021; Domènech-Arúmi 2025). Third, as a technical advantage, the homogeneous kernel gives the decay parameter α a direct interpretation as the rate at which the influence of inequality diminishes with distance, and it is differentiable in α , which permits the analytical characterization of how LII_j responds to the decay parameter. This specification includes $\alpha = 0$ as a special case at which the kernel is constant and

⁴Specifically, setting $w(d_{ij}) = \mathbf{1}\{d_{ij} \leq \bar{d}\}$, where $\mathbf{1}\{\cdot\}$ denotes the indicator function and $\bar{d} > 0$ is a fixed radius, recovers a fixed-radius neighbourhood construction of the type used by Domènech-Arúmi (2025), here applied to MLD, in which all units within $\bar{d}$ receive equal weight, while those beyond $\bar{d}$ receive zero weight.

the LII reduces to the global MLD. Researchers can therefore calibrate α according to their assumptions about the spatial scale of inequality perception, rather than having this decay pattern implicitly determined by the choice of kernel function.

4 Properties of the Local Inequality Index

4.1 Properties independent of the kernel

Even without a specific definition of the kernel function, LII inherits several elementary properties from its GE-family lineage, collected in Proposition 1.

Proposition 1 (Elementary properties). Let j be a focal unit and let the kernel satisfy $\sum_t w(d_{jt})N_t > 0$. Then:

- (i) **Scale invariance.** For any $c > 0$, $LII_j(c\bar{X}_1, \dots, c\bar{X}_n) = LII_j(\bar{X}_1, \dots, \bar{X}_n)$.
- (ii) **Continuity.** LII_j is continuous, and in fact differentiable, in $(\bar{X}_1, \dots, \bar{X}_n)$.
- (iii) **Replication invariance.** Holding the self-distance d_{jj} fixed, if every spatial unit i is replicated together with an identical copy located at the same distance from j as the original (so that the distance structure and each $\bar{X}_i$ are preserved), LII_j is unchanged.

These properties are inherited from LII's GE lineage: spatialization only reweights $S_j(p_i)$ by distance and population, leaving the algebraic structure that gives GE measures these properties untouched. The qualification on replication invariance is necessary, however: without holding d_{jj} fixed, re-applying Equation (4) after replication would shrink d_{jj} relative to the copy, changing LII_j . The proof is in the Appendix. Property (iii) is an algebraic property conditional on fixed d_{jj} ;

it does not imply invariance of the index to splitting a unit or redrawing spatial units, since Equation (4) changes d_{jj} in those cases.

Remark 1 (Anonymity). LII is invariant to joint relabelling of incomes, populations, positions and focal labels, but not to permutations of incomes across fixed locations. This is a consequence of its construction, in which the kernel enters the index itself, rather than a logical implication of Spatiality (Property 2).

The LII also has the following properties, central to its role as an inequality measure.

Proposition 2 (Non-negativity and normalization). $LII_j \geq 0$ for all j , and $LII_j = 0$ if and only if $\bar{X}_i$ is constant across all i with $w(d_{ij})N_i > 0$.

Non-negativity is inherited in the same way as the properties of Proposition 1: $S_j(p_i)$ remains a valid probability weighting, so the Jensen's inequality argument behind $MLD \geq 0$ applies unchanged. The proof is in the Appendix.

However, not every property carries over from GE family. The Pigou-Dalton transfer principle does not hold unconditionally: because the locations of the donor and the recipient enter the index itself, a transfer from a rich but remote unit to a less rich but nearby one can widen the inequality at j 's location even as it narrows the inequality in terms of the national distribution. We therefore state a spatial version of the principle:

Proposition 3 (Spatial transfer principle). For units r and q with $\bar{X}_r > \bar{X}_q$, if $w(d_{jr}) = w(d_{jq}) > 0$, any mean-preserving progressive transfer of an amount $T > 0$ from r to q that does not reverse their ranking strictly decreases LII_j .

Proposition 3 gives a sufficient condition: it covers only donor–recipient pairs that are equally weighted from j . For a general transfer of T from r to q with populations fixed,

$$\frac{dLII_j}{dT} = \frac{1}{\sum_t w(d_{jt})N_t} \left[w(d_{jq}) \left(\frac{1}{\overline{X}_j^w} - \frac{1}{\overline{X}_q} \right) + w(d_{jr}) \left(\frac{1}{\overline{X}_r} - \frac{1}{\overline{X}_j^w} \right) \right],$$

which can have either sign because the transfer also moves $\overline{X}_j^w$.

The next subsection develops the properties of LII under the homogeneous kernel specified in Equation (5), and characterizes how the index varies with the decay parameter α .

4.2 Properties under the homogeneous kernel

The decay parameter α governs how widely the reference environment of each location is drawn.

Proposition 4 (Properties under the homogeneous kernel). Let the kernel function take the homogeneous form, $w(d) = d^{-\alpha}$ with $\alpha \geq 0$, and let the self-distance be set by $d_{jj} = k \min_{i \neq j} d_{ij}$ with $0 < k \leq 1$. Then:

- (i) **Distance invariance.** Rescaling all distances by a common factor $c > 0$, that is $d_{ij} \mapsto c d_{ij}$ for every pair $i \neq j$ and hence $d_{jj} \mapsto c d_{jj}$, leaves LII_j unchanged.
- (ii) **Extreme cases of α .** Let $m_j = \min_{i \neq j} d_{ij}$, let $M_j = \{j\} \cup \{i \neq j : d_{ij} = m_j\}$ collect the focal unit together with its nearest neighbours. Then

(a) if $\alpha = 0$, or more generally if $w(d_{ij}) = \bar{w} > 0$ for all i, j , then

$$LII_j = MLD \quad \text{for every } j;$$

(b) as $\alpha \rightarrow \infty$,

$$LII_j \rightarrow \begin{cases} 0, & k < 1, \\ \sum_{i \in M_j} \frac{N_i}{\sum_{t \in M_j} N_t} \ln \frac{\sum_{t \in M_j} N_t \bar{X}_t}{\left(\sum_{t \in M_j} N_t\right) \bar{X}_i}, & k = 1. \end{cases}$$

(iii) **Response to α .** For any $\alpha \geq 0$, the derivative being one-sided at $\alpha = 0$,

$$\frac{\partial LII_j}{\partial \alpha} = -\text{Cov}_j(D_i^{(j)}, \ln d_{ij}), \quad (6)$$

where

$$D_i^{(j)} \equiv \frac{\bar{X}_i}{\bar{X}_j^w} - \ln \frac{\bar{X}_i}{\bar{X}_j^w} \quad (7)$$

is the local deviation of unit i from the local reference mean, and $E_j[U] = \sum_i S_j(p_i) U_i$ and $\text{Cov}_j(U, V) = E_j[UV] - E_j[U] E_j[V]$ are the mean and the covariance taken under the spatial weights centred on j .

Proposition 5 (Locality and Spatiality). Let the kernel be homogeneous, $w(d) = d^{-\alpha}$ with $\alpha > 0$, and let the self-distance be set by Equation (4). Then LII satisfies Locality (Property 1) and Spatiality (Property 2). At $\alpha = 0$ it satisfies neither.

Proposition 4(i) shows that a common rescaling of all distances leaves the index unchanged, so that LII_j depends on geography only through relative distances; this fails for any kernel fixing an absolute distance scale, such as the Gaussian with

a fixed bandwidth or the fixed-radius indicator.

Proposition 4(ii)(a) nests the global MLD in LII_j at $\alpha = 0$. For (ii)(b), if $k < 1$ the self-distance is strictly smaller than the distance to any other unit, so the weight on j itself dominates entirely, $\overline{X}_j^w \rightarrow \overline{X}_j$, and the index vanishes. If $k = 1$, the focal unit and its nearest neighbours all lie at distance m_j and receive the same kernel weight, so LII_j converges instead to the population-weighted MLD over M_j — a positive value unless every member of M_j shares the same mean income; M_j reduces to a pair whenever the nearest neighbour is unique. Steepening the decay accordingly does not make the index diverge; it contracts the environment that j evaluates.

Proposition 4 traces the full trajectory of LII_j across α , and it is worth pausing on the shape of that trajectory before turning to its proof, given in the Appendix. At $\alpha = 0$, every LII_j collapses to the same value, the global MLD. As α rises, units begin to diverge from the global MLD, moving up or down according to Equation (6), so that a single number splits into as many trajectories. Yet this dispersion does not run away to infinity. As $\alpha \rightarrow \infty$, each LII_j settles again, this time on a value determined entirely by j 's immediate neighbourhood: 0 if j 's nearest neighbour is farther than j 's own self-distance ($k < 1$), or the population-weighted MLD over M_j if they are equidistant ($k = 1$). The index is thus pinned down at both ends, at a single shared value when $\alpha = 0$ and at a strictly local one as $\alpha \rightarrow \infty$, with the whole of the intervening behaviour governed by Equation (6).

Equation (6) shows how α moves the index. Raising α shifts weight from distant units towards nearby ones, and whether this raises or lowers LII_j depends on where, within j 's neighbourhood, the largest deviations from the local reference mean happen to sit. The covariance in Equation (6) answers exactly this: it is negative

when the units deviating most from $\bar{X}_j^w$ are also the closest to j , in which case sharpening the decay concentrates weight where the dispersion already is and LII_j rises with α ; it is positive when the largest deviations lie instead on the periphery while j 's immediate surroundings are comparatively homogeneous, in which case sharpening the decay withdraws weight from that dispersed periphery and LII_j falls.

Because $S_j(p_i)$, and with it $\bar{X}_j^w$ and $D_i^{(j)}$, changes with α , this covariance can itself change sign as α varies, so Equation (6) need not preserve its sign across the full range and LII_j need not move monotonically in α .

5 Empirical illustration

5.1 China's Local Inequality Indexes Evolution

We use data on about 290 Chinese prefecture-level cities, depending on data availability, from 2000 to 2020, drawing per capita GDP and resident population from the Chinese Urban Statistical Yearbooks, with distances computed between locations of principal offices. Using these data, we calculate the LIIs in different years to observe local inequality dynamics, setting $\alpha = 0.5$ and $k = 0.5$, and under alternative values of α to examine the characteristics of LII. Because $\bar{X}_i$ is city-level per capita GDP, the LII here measures income gaps between cities in each city's reference environment.

Figure 1 displays the spatial distribution of LII at five-year intervals from 2000 to 2020, with the colour scale centred on that year's median LII. The national level of LII declines over time, indicating a more equal income distribution. This result

aligns with the conventional MLD estimates, shown in Table 2.

Table 2: Summary Statistics of LII


Year	Cities	Max	Min	Med	MLD
2000	287	0.271	0.204	0.221	0.227
2005	290	0.294	0.209	0.229	0.234
2010	295	0.217	0.165	0.180	0.182
2015	292	0.160	0.128	0.138	0.140
2020	295	0.141	0.112	0.124	0.126

Note: cities can be absent in specific years due to data availability. MLD is the conventional mean log deviation estimate, using population as weight.


The conventional MLD estimate declines from 0.227 in 2000 to 0.126 in 2020, after a slight rise in 2005, and the maximum, minimum and median LII move with it. Notably, the median LII is consistently close to the conventional MLD value in corresponding years. This indicates that LII tracks the conventional MLD’s dynamics closely. This is expected: by Proposition 4(ii)(a) the LII starts from the global MLD at $\alpha = 0$, and under moderate decay most locations remain close to it. The choice of $\alpha = 0.5$ is illustrative; Figure 4 shows how the index responds to other values.

We next turn to the spatial patterns. Given the declining trend of inequality in China, we concentrate on the rank and deviation from the median in each year, instead of using a universal rule across years.


Several patterns emerge from the maps. First, the central and western regions maintain stable LII levels below the national median throughout the sample period. Stability here refers to rank, not level. Their LII values fall over time together with national inequality, but they remain systematically in the lower part of the contemporaneous distribution in every year. The compression of the national




(a) 2000




(b) 2005



(c) 2010



(d) 2015



(e) 2020

Figure 1: Local Inequality Index (LII) across Chinese prefecture-level cities, 2000–2020


distribution therefore does not alter their relative position. However, a low LII does not necessarily indicate a favourable outcome. In the central and western regions, it reflects a spatially homogeneous but low-income environment, in which no city stands out from its neighbourhood, rather than an improvement in the local income distribution.

Second, the Yangtze River Delta (YRD) and Pearl River Delta (PRD) regions remain the clusters of high LII values. In YRD, Shanghai has been among the top two richest cities in China throughout the sample period, yet neither Shanghai nor its surroundings record the highest LII in the country. Table 3 shows why. The YRD is not organised around a single dominant centre: the cities within 100 km of Shanghai are themselves among the most developed in China, and by 2020 their population-weighted mean reaches 88% of Shanghai. A high income level does not by itself produce a dispersed local income environment.


On the other hand, the PRD presents a different configuration. Shenzhen's income is about twice that of the cities within 100 km, whereas the corresponding ratio in the YRD falls from 1.8 to 1.1 over the same period. The steepness of this local gradient sustains the high LIIs in the region. Until 2015 the city with the highest LII in the PRD is Guangzhou rather than Shenzhen. Guangzhou's income falls well short of Shenzhen's while lying close enough for Shenzhen to raise its spatially weighted neighbourhood mean.

The pattern illustrates that LII aggregates the income gradient across the full neighbourhood, with the contribution of each city adjusted by distance. The city with the highest LII is not necessarily the city with the highest, nor the lowest, income within the area.


The Northeast follows a different path. Its LII values start below the national




(a) 2016



(b) 2017



(c) 2018



(d) 2019

Figure 2: Local Inequality Index across northeastern areas: 2016–2019

Table 3: Population-weighted mean per capita GDP by distance band from Shanghai and Shenzhen (yuan)

Year	Yangtze River Delta (Shanghai)				Pearl River Delta (Shenzhen)			
	Centre	<100	100–300	300–1000	Centre	<100	100–300	300–1000
2000	30307	16459	14641	7285	33276	15697	10200	6491
2005	49377	35870	27885	13714	61843	30757	18314	11406
2010	79396	69958	55226	28938	99095	51801	33348	25313
2015	109186	99703	84843	47937	135271	69198	59674	44874
2020	156803	138429	120591	68290	159820	89209	78590	66885

median, rise steadily, and by 2020 the entire region lies above it, reflecting a worsening position of the region within the national income geography. The details of this transformation can be observed from Figure 2.⁵

Table 4: Income dispersion and relative income in Northeast China

Year	N	MLD		Mean income		
		Northeast	National	Northeast	National	Ratio
2010	33	0.1296	0.1824	33568	31873	1.053
2015	34	0.1001	0.1403	44085	52479	0.840
2016	34	0.0875	0.1384	42633	56611	0.753
2017	34	0.0855	0.1376	45184	62864	0.719
2018	28	0.0738	0.1329	51953	68116	0.763
2019	30	0.0924	0.1320	50324	73028	0.689
2020	34	0.0814	0.1257	51751	75523	0.685

Table 4 explains why this transformation occurs. Within the Northeast, income inequality between these cities falls over the period, apart from a rebound in 2019, and faster than the national series. The rise in LII is therefore not driven by disparities among northeastern cities themselves, but by the gap with the rest of the country. The population-weighted mean per capita GDP of the Northeast stands slightly above the national mean in 2010, at a ratio of 1.053, and falls to

⁵Missing observations for certain cities may affect year-to-year comparisons; for example, some cities' LII in 2018 turns from red to blue, then immediately worsens in 2019.




Figure 3: The LII map in 2020, with $\alpha = 4$




Figure 4: Sensitivity of LII to the decay parameter α

0.685 by 2020, with a temporary rise in 2018 when the sample drops to 28 cities. Because the weighting kernel extends beyond regional boundaries, this external gap enters the local income environment of every northeastern city, and their LII values rise even as the region becomes internally more homogeneous. This is a configuration that a single national or regional figure cannot detect: both the national MLD and the regional MLD decline, and only a local measure registers the reversal. We leave the mechanisms behind the divergence to future work.

We close the empirical section with a brief illustration of how α shapes the LII construction itself. Figure 3 recomputes the 2020 map at $\alpha = 4$. The hotspots no longer cover Beijing, Shanghai, and Guangzhou, but surround them instead. Figure 4 traces the mechanism, plotting LII against α for selected cities. Beijing and Shanghai fall toward zero as α grows: since $k < 1$, the self-weight comes to dominate as the kernel steepens, so $\bar{X}_j^w \rightarrow \bar{X}_j$ and their local reference mean converges to their own income. Because these cities have much poorer surrounding cities, a higher α relocates the highest LII values to their immediate periphery instead: Baoding, for instance, rises well above the median, as its reference mean stays anchored to its far richer neighbour, Beijing.

6 Conclusion

This paper proposes the LII to measure inequality at a location, which we characterise by two properties, Locality and Spatiality. By embedding a distance-decay kernel with $\alpha > 0$, the LII satisfies both properties, and position shapes both the local reference mean and the weight each disparity carries. Under a homogeneous kernel, the LII nests the global MLD at $\alpha = 0$. Applied to Chinese prefecture-level

cities, the index reveals patterns that no single-figure index can register: isolated high-inequality agglomerations, and a Northeast where local inequality rises while both national and regional MLD decline. By attaching a value to each location, the LII turns this masked variation into a map of where the income disparities in the local environment are largest. It does so at the resolution of the data themselves: city-level data yield a city-level map, with no need to aggregate cities into provinces.

Several directions remain open for future research. One is to develop a principled selection rule for α which reveals how narrowly a unit's income environment is drawn; an empirical anchor is whether rankings by LII stay stable across plausible choices of α , or diverge in ways that reveal the width of the relevant comparison group. Another follows from the additive decomposability of the GE family. Inequality decomposition in the current literature (Akita 2003; Mu and Li 2026) could then be carried out location by location. The same decomposability allows each LII_j to be split by sector or demographic group, asking not only how unequal a location's environment is but what drives the inequality at this location.

More broadly, the LII is not restricted to geographic space. It applies to any setting in which a distance can be meaningfully defined. Weighting individuals by their distance in age from a focal age, for instance, would yield an age-specific inequality profile, one that reflects how unequally each cohort is positioned relative to those close to it in age rather than to the population as a whole; the same logic extends to social networks, where Stark, Bielawski, and Falniowski (2024) take the comparison group to be the set of units to which a member is directly linked.

Funding: This research received no specific funding.

References

- Takahiro Akita (2003). “Decomposing regional income inequality in China and Indonesia using two-stage nested Theil decomposition method”. *The Annals of Regional Science* 37.1, pp. 55–77. DOI: [10.1007/s001680200107](https://doi.org/10.1007/s001680200107).
- Francesco Andreoli and Eugenio Peluso (2021). “Inference for the neighbourhood inequality index”. *Spatial Economic Analysis* 16.3, pp. 313–332.
- Luc Anselin (1995). “Local Indicators of Spatial Association—LISA”. *Geographical Analysis* 27.2, pp. 93–115.
- Benjamin A Austin, Edward L Glaeser, and Lawrence H Summers (2018). “Jobs for the Heartland: Place-based policies in 21st century America”.
- Ana Condeço-Melhorado, Hande Demirel, Mert Kompil, Elena Navajas, and Panayotis Christidis (2016). “The impact of measuring internal travel distances on selfpotentials and accessibility”. *European Journal of Transport and Infrastructure Research* 16.2, pp. 300–318. DOI: [10.18757/ejtir.2016.16.2.3139](https://doi.org/10.18757/ejtir.2016.16.2.3139).
- Guillermo Cruces, Ricardo Perez-Truglia, and Martin Tetaz (2013). “Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment”. *Journal of public Economics* 98, pp. 100–112.
- Silvia De Nicolò and Maria Rosaria Ferrante (2025). “Mapping Income Inequality Trends Between Spatial and Socio-Demographic Disparities”. *Review of Income and Wealth* 71.2, e70019.
- Gerard Domènech-Arumí (2025). “Neighborhoods, perceived inequality, and preferences for redistribution: Evidence from barcelona”. *Journal of Public Economics* 242, p. 105288.

- Keith Head and Thierry Mayer (2010). “Illusory border effects: distance mis-measurement inflates estimates of home bias in trade”. *The Gravity Model in International Trade: Advances and Applications*. Ed. by Peter A. G. van Bergeijk and Steven Brakman. Cambridge: Cambridge University Press, pp. 165–192.
- Morgan Kelly (2000). “Inequality and crime”. *Review of economics and Statistics* 82.4, pp. 530–539.
- Erzo F. P. Luttmer (2005). “Neighbors as Negatives: Relative Earnings and Well-Being”. *The Quarterly Journal of Economics* 120.3, pp. 963–1002. DOI: [10.1093/qje/120.3.963](https://doi.org/10.1093/qje/120.3.963).
- Miguel A Márquez, Elena Lasarte-Navamuel, and Marcelo Lufin (2018). “Isolating neighborhood components of regional inequality: Illustration for the Spanish case”. *International Regional Science Review* 41.5, pp. 483–509.
- Dongyuan Mu and Junjie Li (2026). “Effect of Public Pensions on Inequality of Opportunity in China”. *Review of Income and Wealth* 72.2, e70069.
- Benjamin J Newman, Christopher D Johnston, and Patrick L Lown (2015). “False consciousness or class awareness? Local income inequality, personal economic position, and belief in American meritocracy”. *American Journal of Political Science* 59.2, pp. 326–340.
- Sergio J Rey and Richard J Smith (2013). “A spatial decomposition of the Gini coefficient”. *Letters in Spatial and Resource Sciences* 6.2, pp. 55–70.
- Andrés Rodríguez-Pose (2018). “The revenge of the places that don’t matter (and what to do about it)”. *Cambridge journal of regions, economy and society* 11.1, pp. 189–209.
- Anthony Shorrocks and Guanghua Wan (2005). “Spatial decomposition of inequality”. *Journal of Economic Geography* 5.1, pp. 59–81.

Oded Stark, Jakub Bielawski, and Fryderyk Falniowski (2024). “Measuring income inequality in social networks”. *The Journal of Economic Inequality* 22, pp. 333–356.

Henri Theil (1967). *Economics and Information Theory*. Amsterdam: North-Holland.

Shlomo Yitzhaki (1979). “Relative deprivation and the Gini coefficient”. *The Quarterly Journal of Economics* 93.2, pp. 321–324.

Appendix

A Proofs of Propositions 1–5

A.1 Proof of Proposition 1

Part (i). The weights $S_j(p_i)$ depend only on distances and populations and are therefore unaffected by a common scaling of incomes. Replacing $\bar{X}_i$ by $c\bar{X}_i$ replaces $\bar{X}_j^w$ by $c\bar{X}_j^w$, so that, by Equation (3),

$$\ln(c\bar{X}_j^w) - \sum_i S_j(p_i) \ln(c\bar{X}_i) = \ln c + \ln \bar{X}_j^w - \ln c - \sum_i S_j(p_i) \ln \bar{X}_i = LII_j,$$

where the cancellation uses $\sum_i S_j(p_i) = 1$.

Part (ii). On the positive orthant $\bar{X}_i > 0$ the map $(\bar{X}_1, \dots, \bar{X}_n) \mapsto LII_j$ is a composition of the smooth function $\ln(\cdot)$ with finite weighted sums, hence differentiable.

Part (iii). Let each unit i be replicated by a copy i' with $N_{i'} = N_i$, $\bar{X}_{i'} = \bar{X}_i$ and $d_{i'j} = d_{ij}$, and hold d_{jj} fixed. Every term $w(d_{ij})N_i$ then appears twice, the

denominator $\sum_t w(d_{jt})N_t$ doubles, and $S_j(p_i) + S_j(p_{i'})$ equals the original $S_j(p_i)$. Both $\bar{X}_j^w = \sum_i S_j(p_i)\bar{X}_i$ and $\sum_i S_j(p_i) \ln \bar{X}_i$ are therefore unchanged, and so is LII_j by Equation (3). If instead Equation (4) is re-applied after replication, the copy j' lies at distance d_{jj} from j and the rule returns $k d_{jj}$ in its place; the unnormalized weight attached to $\bar{X}_j$ becomes $N_j[w(k d_{jj}) + w(d_{jj})]$ while that of every other unit merely doubles, so the weight of the focal unit relative to any other is multiplied by $[1 + w(k d_{jj})/w(d_{jj})]/2 > 1$ and LII_j changes. Under the homogeneous kernel this factor is $(1 + k^{-\alpha})/2$. $\square$

A.2 Proof of Proposition 2

Since the weights $S_j(p_i)$ form a probability distribution over the units, the LII in Equation (3) can be rewritten as:

$$LII_j = \ln \bar{X}_j^w - \sum_i S_j(p_i) \ln \bar{X}_i$$

Since $\ln(\cdot)$ is strictly concave, Jensen's inequality yields:

$$\sum_i S_j(p_i) \ln \bar{X}_i \leq \ln \left(\sum_i S_j(p_i) \bar{X}_i \right) = \ln \bar{X}_j^w$$

Therefore $LII_j \geq 0$, with equality if and only if $\bar{X}_i$ is constant across the units carrying positive weight: strict concavity makes the inequality strict as soon as two of them differ in income, while units with $S_j(p_i) = 0$, whether because $w(d_{ij}) = 0$ or because $N_i = 0$, are unrestricted. $\square$

A.3 Proof of Proposition 3

Under the transfer, in which an amount $T > 0$ of income is moved from r to q while populations are held fixed, $d\bar{X}_r = -T/N_r$ and $d\bar{X}_q = T/N_q$. The weights $S_j(p_i)$ depend on distances and populations only, so the transfer leaves them unaltered and only the incomes need be differentiated. Since $S_j(p_i) = w(d_{ij})N_i / \sum_t w(d_{jt})N_t$, the change in the spatially weighted mean is:

$$d\bar{X}_j^w = S_j(p_q) \frac{T}{N_q} - S_j(p_r) \frac{T}{N_r} = \frac{T[w(d_{jq}) - w(d_{jr})]}{\sum_t w(d_{jt})N_t} = 0$$

since $w(d_{jq}) = w(d_{jr})$. With $d\bar{X}_j^w = 0$, differentiating $LII_j = \sum_i S_j(p_i) \ln(\bar{X}_j^w / \bar{X}_i)$ yields:

$$dLII_j = -S_j(p_q) \frac{T/N_q}{\bar{X}_q} + S_j(p_r) \frac{T/N_r}{\bar{X}_r} = \frac{w(d_{jq})T}{\sum_t w(d_{jt})N_t} \left(\frac{1}{\bar{X}_r} - \frac{1}{\bar{X}_q} \right) < 0$$

since $\bar{X}_r > \bar{X}_q$ and $w(d_{jq}) > 0$. For a transfer of finite size the same expression is integrated over the path, and its integrand retains the sign of $1/\bar{X}_r - 1/\bar{X}_q$ throughout precisely when the ranking of the two units is not reversed, which is the condition imposed in the statement. Without it the conclusion fails: once the transfer carries $\bar{X}_q$ above $\bar{X}_r$, further transfers increase LII_j , and a sufficiently large transfer leaves the index above its initial value. Dropping the hypothesis $w(d_{jr}) = w(d_{jq})$ instead yields the general expression given after the proposition in the text. $\square$

A.4 Proof of Proposition 4

For (i), rescaling every distance by c , including d_{jj} , multiplies every kernel weight by $c^{-\alpha}$, which cancels in $S_j(p_i)$. For (ii)(a), a constant kernel gives $S_j(p_i) = N_i/N = p_i$ for every j , so LII_j reduces to Equation (1).

Part (ii)(b). Note that $S_j(p_i)/S_j(p_j) = (N_i/N_j)(d_{jj}/d_{ij})^\alpha$. If $k < 1$ then $d_{jj} \leq k d_{ij}$ for every $i \neq j$, so each such ratio is bounded by $(N_i/N_j)k^\alpha \rightarrow 0$; with n finite, $S_j(p_j) \rightarrow 1$, $\overline{X}_j^w \rightarrow \overline{X}_j$ and $LII_j \rightarrow 0$ by Equation (3). If $k = 1$, then $d_{jj} = m_j$, so that $d_{ij} = d_{jj}$ for every $i \in M_j$ while $d_{ij} > d_{jj}$ for all remaining i . The weight of every unit outside M_j therefore vanishes and

$$S_j(p_i) \longrightarrow \frac{N_i}{\sum_{t \in M_j} N_t}, \quad i \in M_j.$$

The kernel weights of the members of M_j coincide, but their normalized weights remain proportional to population, so that $\overline{X}_j^w \rightarrow \sum_{t \in M_j} N_t \overline{X}_t / \sum_{t \in M_j} N_t$ and, by Equation (3), LII_j converges to the population-weighted MLD over M_j given in the statement. This limit is strictly positive unless $\overline{X}_i$ is the same for every $i \in M_j$.

Part (iii). Since $\overline{X}_i$ and d_{ij} do not depend on α , the parameter enters LII_j only through the weights $S_j(p_i)$. Writing the kernel in exponential form, $w(d_{ij}) = e^{-\alpha \ln d_{ij}}$, and taking logarithms,

$$\ln S_j(p_i) = -\alpha \ln d_{ij} + \ln N_i - \ln \sum_t w(d_{jt}) N_t.$$

The derivative of the last term in α is $-E_j[\ln d_{ij}]$, so that

$$\frac{\partial S_j(p_i)}{\partial \alpha} = -S_j(p_i) \left[\ln d_{ij} - E_j[\ln d_{ij}] \right].$$

Since $\sum_i S_j(p_i) = 1$ for every α , these derivatives sum to zero, confirming that the weights remain a valid probability distribution along the path.

Recall the Jensen-gap representation $LII_j = \ln \bar{X}_j^w - \sum_i S_j(p_i) \ln \bar{X}_i \equiv L_1 - L_2$, where $\bar{X}_j^w = \sum_i S_j(p_i) \bar{X}_i$. Differentiating L_2 using the weight derivative above,

$$\frac{\partial L_2}{\partial \alpha} = \sum_i \frac{\partial S_j(p_i)}{\partial \alpha} \ln \bar{X}_i = - \sum_i S_j(p_i) [\ln d_{ij} - E_j[\ln d_{ij}]] \ln \bar{X}_i = -\text{Cov}_j(\ln \bar{X}_i, \ln d_{ij}),$$

where the last equality is the definition of Cov_j . The identical argument applied to $\bar{X}_j^w = \sum_i S_j(p_i) \bar{X}_i$ gives

$$\frac{\partial \bar{X}_j^w}{\partial \alpha} = -\text{Cov}_j(\bar{X}_i, \ln d_{ij}),$$

$$\frac{\partial L_1}{\partial \alpha} = \frac{1}{\bar{X}_j^w} \frac{\partial \bar{X}_j^w}{\partial \alpha} = -\frac{\text{Cov}_j(\bar{X}_i, \ln d_{ij})}{\bar{X}_j^w}.$$

Subtracting,

$$\frac{\partial LII_j}{\partial \alpha} = \frac{\partial L_1}{\partial \alpha} - \frac{\partial L_2}{\partial \alpha} = -\frac{\text{Cov}_j(\bar{X}_i, \ln d_{ij})}{\bar{X}_j^w} + \text{Cov}_j(\ln \bar{X}_i, \ln d_{ij}).$$

Since $\bar{X}_j^w$ is constant across the summation index i , linearity of $\text{Cov}_j(\cdot, \cdot)$ in its first argument allows the first term to be absorbed as

$$-\frac{\text{Cov}_j(\bar{X}_i, \ln d_{ij})}{\bar{X}_j^w} = -\text{Cov}_j\left(\frac{\bar{X}_i}{\bar{X}_j^w}, \ln d_{ij}\right),$$

which combines with the second term into

$$\frac{\partial LII_j}{\partial \alpha} = \text{Cov}_j\left(\ln \bar{X}_i - \frac{\bar{X}_i}{\bar{X}_j^w}, \ln d_{ij}\right).$$

The first argument differs from $-D_i^{(j)}$, as defined in Equation (7), only by the additive term $\ln \bar{X}_j^w$, which does not vary with i and therefore leaves the covariance unchanged; the right-hand side is thus $-\text{Cov}_j(D_i^{(j)}, \ln d_{ij})$, which is Equation (6). At $\alpha = 0$ the derivative is understood as a right-derivative, the expression above being continuous in α on $[0, \infty)$.

A.5 Proof of Proposition 5

Locality. Take three units with equal populations, $\bar{X}_1 = \bar{X}_2 = 1$, $\bar{X}_3 = r \neq 1$, $d_{12} = \varepsilon$ and $d_{13} = d_{23} = D$ with $\varepsilon < D$. At unit 3 the nearest distance is D , so $d_{33} = kD$ and the weights are proportional to $(1, 1, k^{-\alpha})$ whatever ε ; since incomes are not all equal and all weights are positive, LII_3 is a positive constant. At unit 1, $d_{11} = k\varepsilon$ and the weights are proportional to $((k\varepsilon)^{-\alpha}, \varepsilon^{-\alpha}, D^{-\alpha})$, so the share of unit 3 tends to zero as $\varepsilon/D \rightarrow 0$. Units 1 and 2 share the same income, hence $LII_1 \rightarrow 0$, and $LII_1 \neq LII_3$ for ε/D small enough.

Spatiality. Fix a configuration in general position with incomes not all equal, a unit i and a neighbourhood V of its position. Since i is not the nearest neighbour of every other unit, there is a unit $j \neq i$ whose unique nearest neighbour t differs from i , so that $d_{jt} < d_{ji}$. By continuity there is an open ball $U \subset V$ around the position of i on which $d_{ji} > d_{jt}$ still holds and which contains the position of no other unit. On U , $d_{jj} = k d_{jt}$ is unchanged and the only kernel weight at j that varies is that of i . Write $s = S_j(p_i)$, and let μ_0 and L_0 denote the mean and the mean of logarithms of the incomes of the other units under their normalised weights at j , which do not vary on U . Then

$$LII_j = \ln((1-s)\mu_0 + s\bar{X}_i) - (1-s)L_0 - s \ln \bar{X}_i.$$

Because $\alpha > 0$, s is strictly decreasing in d_{ij} , and d_{ij} takes an interval of values on U , so s ranges over an interval. If LII_j were constant on U , its second derivative in s , $-(\bar{X}_i - \mu_0)^2 / ((1-s)\mu_0 + s\bar{X}_i)^2$, would vanish, giving $\bar{X}_i = \mu_0$; its first derivative would then reduce to $L_0 - \ln \mu_0 = 0$. All other units carry positive weight, so by strict concavity of $\ln(\cdot)$ this requires them to share a common income, which equals $\mu_0 = \bar{X}_i$; all incomes would then be equal, a contradiction. Hence some position in $U \subset V$ changes LII_j .

The case $\alpha = 0$. By Proposition 4(ii)(a), $LII_j = MLD$ for every j and every configuration, so the index neither differs across locations nor depends on positions.

□