

Confounded Local Inference: Extending Local Moran Statistics to Handle Confounding

Levi John Wolf *

School of Geographical Sciences
University of Bristol
levi.john.wolf@bristol.ac.uk

October 27, 2023

Abstract

Local statistical analysis has long been of interest to social and environmental scientists who analyze geographic data. Research into local spatial statistics experienced a step-change in the mid-1990s, which provided a large class of local statistical methods and models. The local Moran statistic is one commonly used local indicator of spatial association, able to detect both areas of similarity and observations that are very dissimilar from their surroundings. From this, many further local statistics have been developed to characterize spatial clusters and outliers. However, these statistics have seen limited adoption because they do not sufficiently model the relationships involved in *confounded spatial data*, where the analyst seeks to understand the local spatial structure of a given outcome variable that is influenced by one or more additional factors. Recent innovations used to do *joint* multivariate local analysis also do not model this kind of *conditional* local structure in data. This paper provides tools to rigorously characterize confounded local inference and provide a new and different class of *multivariate conditional local Moran statistics* that can account for confounding. To do this, we will return to the Moran Scatterplot as the critical tool for local Moran-style covariance statistics. Extending this concept, a new method is available directly from a “Moran-form” multiple regression. We show the empirical and theoretical properties of this statistic, show how some existing heuristic approaches arise naturally from this framework, and show how the use of conditional inference can change interpretations in an empirical analysis of rent and housing stock in a rapidly changing neighborhood.

Introduction

Making maps of local indicators of spatial association and interpreting their meaning is a common practice among applied geographers. At the frontiers of spatial methodology, local analysis

*Thank you to Profs. Sergio Rey, Luc Anselin, and Alex Marsh for helping me develop the idea beyond the initial proof of concept and giving feedback on previous drafts. This work is part of an ESRC funded Impact Acceleration project through the University of Bristol, *Bristol City Council Living Rent Commission: Understanding the Local Structure of Affordable Housing*, and further by PolicyBristol from the Research England QR Policy Support Fund (QR PSF) 2022-24.

has long been recognized as important (Fotheringham 1993; Fotheringham and Brunsdon 1999), with multilevel models (Jones 1991), local indicators of spatial association (LISA) (Anselin 1995), scan statistics (Rogerson 2022) and so-called “partial-map” models (Wolf, Oshan, and Fotheringham 2018) providing many useful frameworks for analyzing distinct local structures in data. One distinction in this literature is often between “exploratory” and “confirmatory” analysis, but all of these local analysis frameworks, at their core, have rigorous statistically-justified frameworks underpinning the way they work.

In particular, some local statistics are used to identify spatial outliers — areas that are markedly dissimilar from their surroundings — as well as spatial clusters — areas of strong similarity. In this context, local statistics have been largely focused on univariate analyses (Getis and Ord 1992; Anselin 1995; Getis and Ord 1996), arising from attempts to rigorously characterize the structure of visual patterns analysts see in univariate choropleth maps. However, we live in a “multivariate world” (Anselin and Li 2020), so recent work has focused intensely on trying to define and analyze multivariate spatial outliers and/or clusters (Wartenberg 1985; Lee 2001; Dray and Jombart 2011; Harris et al. 2014; Filzmoser, Ruiz-Gazen, and Thomas-Agnan 2014; Eckardt and Mateu 2021). Often, these multivariate measures of local association are deployed when more than one covariate is relevant to understanding the local spatial structure of a given dataset.

From here, it becomes important to think about two “senses” of *multivariate* that are used in the literature. Some are interested in *joint* spatial association—the spatial patterns of co-occurrence of values—while others refer to *conditional* spatial association—the spatial pattern of one variable *given* its association with others. To illustrate, consider an example for urban house rents to which we will return. A *joint* spatial outlier might be an expensive house on a large lot located in a neighborhood with relatively smaller and cheaper houses. Thus, when we work with *joint* multivariate local statistics, we characterize how distinctive a site is relative to its surroundings using multiple variables of comparison all at the same time. Most methods developed in recent years are of this type, and Anselin and Li (2020) provides an excellent discussion of this practice.

A *conditional* spatial outlier, on the other hand, would reflect a house that is unusually expensive for its area *given its lot size*—a site that is unusual for its area along a specific trait after accounting for its other traits. In reality, it is nearly always necessary for spatial analysts to understand the spatial patterning in a given outcome *taking into account* other causes, correlates, or associates of that outcome. This is most useful in the case of *confounding*, when some other variable influences the outcome under study both at the site of measurement *and* at surrounding sites. This case for a *conditional* analysis often pushes people towards full spatial regression models (Anselin and Griffith 1988), encouraging a shift from “local” to “global” analysis as a result. To try to provide local statistics useful for this purpose, both Anselin, Syabri, and Smirnov (2002) and Lee (2001) revisit Wartenberg (1985) in order to build a measure of “multivariate similarity conditioned on [sites] geographic proximity” (p. 280 Wartenberg 1985). Lee (2001) provides a strong critique of Wartenberg (1985)’s statistic and Eckardt and Mateu (2021) extend upon this critique further, which I discuss below. Regardless, proposed solutions still do not adequately model the associations

between an outcome and potential *confounders* for that outcome. Since uni- and bi-variate local Moran statistics are extensively used in exploratory spatial data analysis, it stands to reason that a fully multivariate local Moran statistic might also be used thoroughly.

Therefore, I develop this statistic. To do so, I first provide a basis for discussion by examining the conceptual underpinnings of Moran statistics. Then, I discuss one common approach for structuring bivariate local Moran statistics (Wartenberg 1985). In doing so, I show how these are incomplete models of the full structure of spatial associations in geographic data. Starting again from first principles, I focus on a multivariate version of a *Moran-form* regression, which yields a multivariate *conditional* local Moran statistic capable of telling us whether an outcome at a given site is unusual *after accounting for* other existing explanatory variables. I show that this method is fundamentally related to another common (but heretofore *ad hoc*) strategy for handling confounding: univariate local analysis of the residuals from an auxiliary regression. Formal insights are drawn about the structure of both estimators by breaking them down into their components, and I show how these components arise from the network of associations between different variables. Finally, I demonstrate their behavior in an analysis of rent in a rapidly-changing neighborhood in Bristol, a metropolitan area of over half a million people in the fastest-growing region of the United Kingdom. I pay specific attention to how multivariable analysis may change interpretations or understandings obtained from univariate local analysis.

Existing Local Covariance Statistics

Local statistics seek to characterize the relationship between each site and its surroundings for a given variable of interest, y . In this case, one of the most thoroughly-used statistics, the local Moran statistic (Anselin 1995), uses the covariance between the specific site's value, y_i , and a measurement of the surrounding site values. These covariance-based Moran statistics are useful because they characterize *both* spatial outliers, observations that are quite different from their surroundings, as well as spatial clusters, where observations are very similar to one another. To extend these univariate statistics into a multivariate analysis that can describe the co-variation between the site and its surroundings, we need three core properties. First, we need the statistic to measure the relationship between a specific outcome y_i at a given site and its surroundings accounting for potential *confounding* information, \mathbf{X} , that may influence the outcome. That is, we want it to be *conditional*. Second, we want this local measure to be one of *co-variation* between two vectors, allowing for the same Moran-type analysis of local outliers and clusters. Third, the statistic must be *local*, in that it measures the co-variation between y (or x) and surrounding y values at each site.

No *conditional local* statistic of *co-variation* currently exists. Methods to obtain direct and indirect effects from global spatial autoregressive models can provide estimates of local association, but they do not directly summarize the co-variation between different y and x variables. Existing local bivariate Moran's I statistics are *local co-variation* statistics, but they do not fully *condition* on \mathbf{X} . Existing auxiliary regression strategies, on the other hand, may *over-correct* for indirect asso-

ciations present in the data. Therefore, we propose a direct extension of the local Moran’s I_i , the *local partially-conditional Moran’s I statistic*, that takes its main point of departure from the Moran Scatterplot. Its corresponding “global” estimate is the slope relating y to $\mathbf{W}y$ obtained from a Moran-form multiple regression estimated using Ordinary Least Squares with \mathbf{X} included in the model. As a consequence, it has a few clear mathematical relationships with existing approaches: it is a correlation-weighted average of existing univariate and bivariate local statistics, and is structurally related to using an auxiliary regression to remove \mathbf{X} from y . However, it is distinct from both of these approaches in important ways. Below, we provide context on the existing local Moran statistics, presenting existing and novel critiques of their usefulness for *confounded local inference*, and then show how this new conditional local statistic improves upon these existing approaches.

The Moran Scatterplot

As a starting framework, Anselin (1996)’s explanation of Moran’s I and local Moran’s I_i in terms of the Moran Scatterplot is exceptionally clear. This scatterplot graphically presents the relationship between the value of y at site i and its “spatial lag,” a summary of other y values near site i . The I is a regression coefficient, so can be thought of as the covariance of y with its spatial lag divided by the variance of y itself. Usually, we adopt some kind of rule to govern which sites are considered “near” one another, and measure them with a weight between zero and one. We usually constrain these weights such that they add up to one for every site, so that the spatial lag of y can be interpreted as a kind of spatially-weighted average of the values near site i . We record the weights relating i to all other sites j in a spatial weights matrix, \mathbf{W} , meaning that the spatial lag of y becomes $\mathbf{W}y$.¹ Thus, the covariation of y with $\mathbf{W}y$ is the main interest for local Moran-type statistics, and the scatterplot relating y to $\mathbf{W}y$ is called the “Moran Scatterplot.” Covariance-based Moran statistics are useful in a large part because we can represent and derive them quickly from this intuitive scatterplot.²

An example of the relationships between the local & global variants of the Moran statistics and their map pattern is shown in Figure 1. This figure shows a Moran scatterplot for rent values in East Bristol from a dataset on the rental listings from a large property market data aggregator, provided by the Urban Big Data Center. Here, y is the weekly rent charged for a given property, so $\mathbf{W}y$ is the typical rent in the area near the property. The dashed lines depict the means of y and $\mathbf{W}y$, respectively, while the slope of the orange regression line is the global Moran’s I in this case. This means that for a z-standardized variable y , we can state the Moran-form regression:

$$\mathbf{W}y = Iy + \epsilon \quad \epsilon \sim \mathcal{N}(0, \sigma^2) \quad (1)$$

From this, we can obtain I , the Moran’s I statistic, by ignoring the fact that y is simultaneous with

¹For the following discussion, we only consider row-standardized weights matrices for simplicity.

²Other relationships besides covariance that are used in other statistics (e.g. distances in the local Geary’s C_i statistic or sums in the Getis-Ord G_i^* statistic) provide different interpretations.

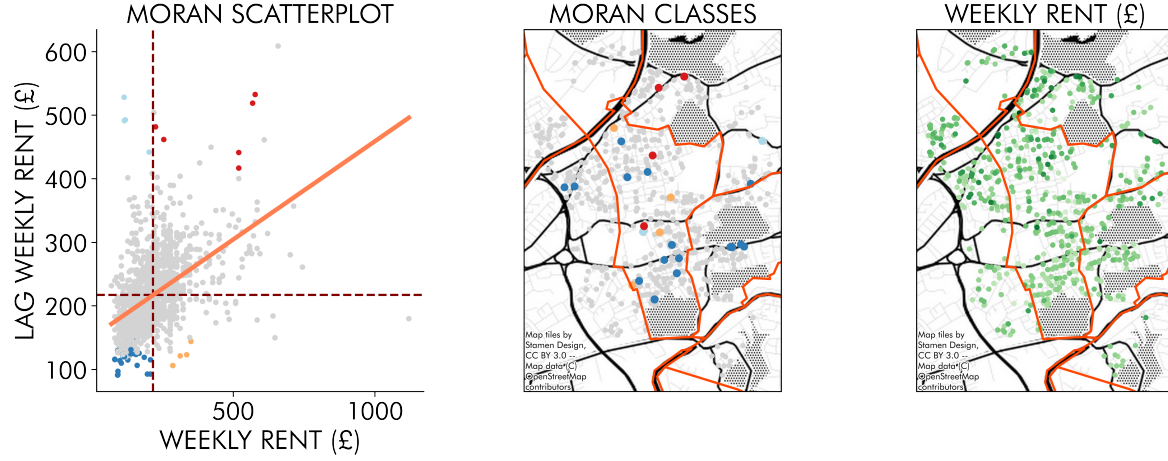


Figure 1: A typical Moran scatterplot (Anselin 1996), showing the values with significant LISA statistics in a non-grey color. Here, the colors reflect the quadrant classifications used in GeoDa (Anselin and Rey 2014), blues represent places where the focal observation (y_i) is below mean and red places where the focal is above mean. The lighter colors reflect outliers, where the relationship between an observation and its surroundings is discordant, and the darker colors reflect concordant areas (clusters).

$\mathbf{W}y$ and using the ordinary least squares estimator:

$$\hat{I} = (y'y)^{-1} (y'\mathbf{W}y) \quad (2)$$

This is the classic Moran's I statistic used throughout exploratory data analyses in geography.

The “localization” strategy for local Moran Statistics

To construct a local indicator of spatial association, Anselin (1995) partitions this estimator into a vector of contributions from each site. Traditionally, local I_i statistics are presented in element summation form over the real numbers, so $\{y_i, x_i, w_{ij}\} \in \mathbb{R}$. However, a vector statement for local statistics will make it simple to show their relationship to the new multivariate conditional local Moran, so I opt for them here. The partitioned vector-form estimator replaces the inner product $y'\mathbf{W}y$ with an elementwise product:

$$\hat{\mathbf{I}} = (y'y)^{-1} (y \circ \mathbf{W}y) = \frac{1}{N} \begin{bmatrix} y_1 \sum_j^N w_{1j} y_j \\ y_2 \sum_j^N w_{2j} y_j \\ \vdots \\ y_n \sum_j^N w_{nj} y_j \end{bmatrix} \quad (3)$$

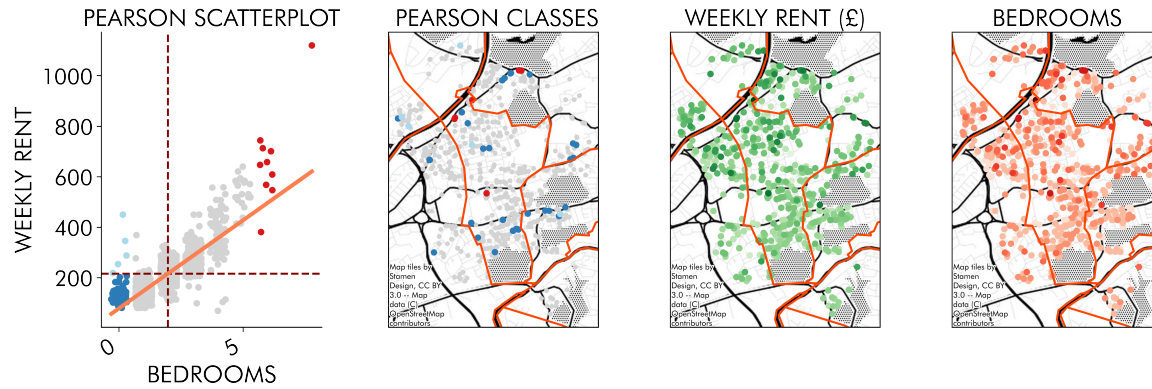


Figure 2: The (a)spatial distribution of rents and rental sizes in the neighborhood of Easton, in Bristol, England. This statistical approach, due to Lee (2001), uses the typical local Moran approach with Pearson correlation, so that “high-high” observations are unusually expensive large rental properties while “low-high” listings are unusually expensive *but small* rental properties. The Easton neighborhood has seen substantial rental market regulation by local government in order to stem the displacement of black and minority ethnic populations that started to occur during 2010s gentrification.

As a consequence of changing from an inner product to an elementwise product, the sum of the local statistics is proportional to the global statistic. This allows us to interpret the local Moran in a very simple fashion. For example, many negative local values add up to a negative global value, since many outliers come together to make a negatively-autocorrelated map overall. Another way to think of local I_i statistics is as a kind of outlier statistic about the Moran Scatterplot (Anselin 1996): those areas that are strongly different from the general relationship between y and W_y are discovered by the local Moran statistics (Anselin 1995). Although, the identified sites will often not reflect a consistent Euclidean distance from the regression line, since the statistics are map- and locality-dependent. This is clear from visual inspection of the scatterplot in Figure 1, where many of the non-grey dots (representing spatially-anomalous observations) are closer to the regression line than some other grey dots. Thus, spatial outliers or clusters may be more than visual outliers in the Moran scatterplot, because their significance depends on the local structure of the map, not necessarily the global trend.

Critically, however, this only displays the univariate relationship between y and W_y . This structure cannot account for the possibility that some third variate, x , may affect the spatial distribution of y . If we were to include x , apparent autocorrelation in y may instead be more simply explained in terms of x . In this case, we can see that there is a very strong (.84) correlation between the rent charged for a property and the number of bedrooms that the property has. This is shown in Figure 2, where the same Moran-type presentation is used to show the Pearson correlation between rent and the number of bedrooms being rented. There, we see that studio apartments (typically single rooms with integrated sleeping-living spaces) shown in dark blue are generally more expensive than would be predicted according to the relationship between property size and price. But, shown in a lighter blue, we can see some studios are distinctively expensive despite

their small size. Thus, it's likely that rent levels are set largely in accordance to the size of the property being let, and then some idiosyncratic factors at the flat- or area-level may then influence the final rent level for a given property. To get at this spatial structure, we need a local measure of the covariance between y and Wy accounting for x that we can use like the typical univariate local Moran's I_i . The next few sections explore a few useful (but ultimately unsuitable) existing methods to do this.

Simultaneous Autoregressive Models

When spatial analysis moves from one to many variables, the Moran-form regression is often abandoned in favor of simultaneous autoregressive (SAR) model (Anselin and Griffith 1988). These models are useful because they (correctly) specify the idea of a process spillover: whereas the Moran scatterplot and its attendant regression treat Wy as just another independent variable, simultaneous autoregressive models treat y as if it arises as a result of the influence of nearby y . Put simply, y is the outcome we observe directly and we can only obtain Wy from it, so we should probably use y as an outcome with W treated as some extra information that helps us predict. This is exactly what simultaneous autoregressive models do, but the Moran-type statistics common in spatial analysis put this “backwards.”

In the analogous spatial econometric model to Equation 1, the spatial autocorrelation in y is estimated along with P additional covariates in some design matrix X . This “flips” the specification from that in Equation 1 so that y is now on the left and Wy is now with the “covariates:”

$$y = \rho Wy + X\beta + \epsilon \quad \epsilon \sim \mathcal{N}(0, \sigma^2) \quad (4)$$

In this form, we can include any x variables we think might be driving the outcome, which means we can account for confounding. Further, estimates from this model reflect estimates of how strongly y is related both to additional information in X and how it is related to itself over the spatial structure in W . Here, ρ is the parameter governing spatial autoregressive structure, β (incompletely³) models the marginal impact of X on y , and ϵ is a typical normally-distributed independent error term.

Unfortunately, local Moran-style statistics are not immediately available from these models. First, we cannot apply the same inner-to-elementwise trick as before because the estimate of spatial structure (ρ) is not separable from the estimate of the “exogenous” information in X :

$$\begin{aligned} y - \rho Wy &= X\beta + \epsilon \\ y(I - \rho W) &= X\beta + \epsilon \\ y &= (I - \rho W)^{-1}(X\beta + \epsilon) \end{aligned} \quad (5)$$

³See, for example, James P. LeSage and Pace (2014)'s argument about the distinction between β and the marginal effect in spatial models.

From this, we can see that the matrix that describes spatial co-variation, $I - \rho\mathbf{W}$, applies across all $\mathbf{X}\beta$ and is not additively separable from $\mathbf{X}\beta$. Any inner-to-elementwise trick we apply in this specification will result in an *interaction* between the spatial filter and the exogenous data.

To try to provide something analagous, James P LeSage and Pace (2009) provides a way to separate the “direct” effect that a variable has on outcomes at the same site from the “indirect” spatial spillovers arising from nearby sites. These come from the on- and off-diagonal elements of $I - \rho\mathbf{W}$, respectively, and apply to each site separately. They measure the effect of an observation on itself or on its neighbors given a change in X , and change depending on the spatial influence of the observation. While these direct/indirect effect estimates are extremely helpful for geographers to reason critically about marginal effects in spatial models, they cannot stand in as a multivariate local covariance statistic because direct/indirect effects apply across all variables equally, while the unique effect of X_j is captured directly in β_j and held constant across sites. Thus, there is no difference in the “indirect” effect of one covariate versus another, nor of that between y and the other potentially-confounding variates \mathbf{X} . This means that these effect estimates cannot even replace common bivariate Moran statistics discussed in the next sections that characterize the spatial covariation of x with $\mathbf{W}y$. They simply obtain from a too-different specification than the classical Moran-style approaches.

Wartenberg (1985)’s Bivariate Moran’s I_i

Returning to the Moran-type statistics, one long-used bivariate local statistic adapts Wartenberg (1985)’s estimator, applying the idea of a scattermatrix to Moran-style comparisons between variables and their spatial lags. For two z-standardized variables y and x , the local bi-variate Moran statistic takes a very similar form to the typical univariate local Moran’s I_i , but this time comparing z_x to z_y :

$$\hat{I}_{xy} = x \circ \mathbf{W}y \quad (6)$$

Further, the order of this grouping matters in most analyses,⁴ so it is common to see the full set of pairwise statistics relating (y, x) each to $\mathbf{W}y$ and $\mathbf{W}x$.

In terms of the Moran scatterplot, the Wartenberg (1985) estimator can be visualized as the “scattermatrix” of pairwise relationships, with each column constituting a different “focal” variate, z_j , and each row constituting a different “neighborhood” variate, $\mathbf{W}z_k$. This is illustrated in Figure 3, where each subplot corresponds to a univariate or bivariate local Moran statistic. The diagonal of this scatter matrix contains the usual univariate local Moran statistic, since it shows the relationship between each covariate and its own lag, and the off-diagonal elements correspond to bivariate local Moran statistics.

Unfortunately, each Wartenberg (1985) estimator only considers the pairwise relationship between one covariate and one lagged covariate at a single time. This means that it ignores the co-variation between $\mathbf{W}y$ and y when examining that for x and $\mathbf{W}y$. Conceptually, this is confusing:

⁴ \mathbf{W} is generally not symmetric when it is row standardized, so $x \circ \mathbf{W}y \neq y \circ \mathbf{W}x$.

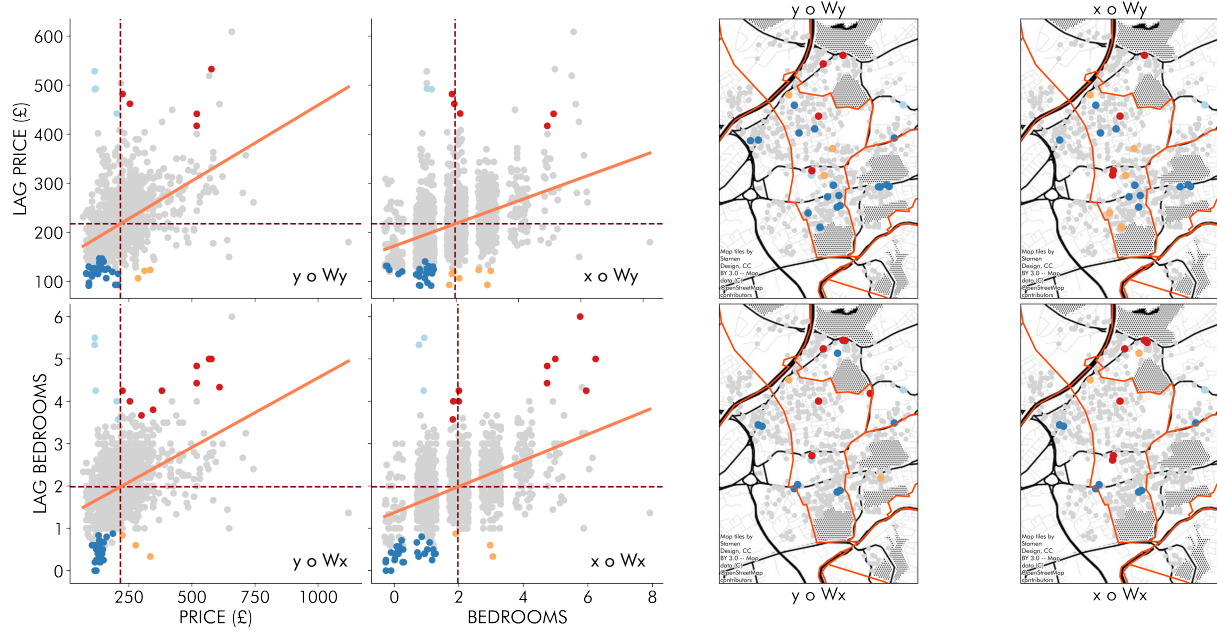


Figure 3: Scattermatrix and bivariate local Moran maps demonstrating the conceptual structure of the Wartenberg bivariate Moran estimator for the covariates shown in Figure 2.

Wy arises simultaneously with y , so the two are mutually-constitutive. If y is observed, so too is Wy ; if y is related to x , it is highly likely that Wy is, too. In terms of our rent and property size example, this estimator would compare each property's size to the average rent in its area, the use of which is not immediately clear. Thus, the conceptualization behind this estimator is useful for specific hypotheses about spatial cross-correlation, but it cannot control for confounding.

Lee (2001)'s Bivariate Statistic

Lee (2001) notes this issue, stating that this means the Wartenberg estimator is “conceptually untenable” (p. 377) and should not be used for analysis. To resolve the issue, an explicit link between the bivariate-form estimator and correlation terms relating two variates is obtained from a decomposition of the classic Moran coefficient into aspatial correlations between variates and *spatial smoothing scalars* that describe how spatial structure increases or reduces the co-variation between them. In matrix form, for a row-standardized weights matrix like that considered here, the Lee (2001)'s bivariate estimator is a *doubly-lagged* version of the Wartenberg (1985) estimator:

$$\mathbf{L} = \mathbf{W}z_x \circ \mathbf{W}z_y \quad (7)$$

Like the Wartenberg estimator, this also results in a matrix of p association coefficients, but the coefficients are now symmetric.

The spatial smoothing scalars and the relationship Lee (2001) establishes between Moran I and Pearson correlation are useful and insightful. But, the L_{xy} statistic shares some of the same prob-

lems as previous bivariate statistics. In the sense of our rent example, the L_{xy} statistic allows us to examine the relationship between local averages of rent and house size, correcting for the mismatch in the bivariate Moran statistic. But, it is still a pairwise measure of association: additional covariates z will not affect an estimated relationship between x and y . The relationship between x and y is completely modelled by a single measure of correlation between some factor (generalized now for a choice of x, y , or $\mathbf{W}x$) and the spatial pattern of y , measured by $\mathbf{W}y$. Further, there is no accounting for the relationship between y and x directly in Eq. 7, only their lagged counterparts. Thus, the measure remains slightly unsatisfactory.

Auxiliary Regression

To step towards a comparison where all relevant pairs of relationships are considered simultaneously, one reasonable approach encountered in the literature is to first remove the contribution that \mathbf{X} makes to y entirely. This is done by assuming any relationship between $\mathbf{W}y$ and \mathbf{X} is a product of that between y and \mathbf{X} alone. A very clear statement of this family of approaches is offered by Eckardt and Mateu (2021) for a general all-pairs spatial cross-correlation statistic. For the purposes of studying spatial confounding, we would use this auxiliary regression strategy to remove the influence of z -standardized exogenous data, \mathbf{X} , from the *primary* z -standardized outcome under study, y . This provides a residual, e , arising from predicting y using \mathbf{X} , which we will then use for analysis in place of y :

$$\hat{\mathbf{I}}_{x \rightarrow y} = e \circ \mathbf{W}e(e'e)^{-1} \quad (8)$$

$$e = y - \hat{y} = y - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'y$$

As an example, we can obtain the following estimator⁵ that describes an auxiliary local Moran statistic using one x variate:

$$\hat{\mathbf{I}}_{x \rightarrow y} = \left[y \circ \mathbf{W}y - \rho_{xy}x \circ \mathbf{W}y - \rho_{xy}y \circ \mathbf{W}x + \rho_{xy}^2 x \circ \mathbf{W}x \right] \frac{(N-1)}{N(1-\rho_{xy}^2)} \quad (9)$$

Thus, we can compute a local Moran statistic using this expression directly, or build e first and use standard univariate local Moran statistics on the result.

This auxiliary regression strategy provides us a unique way of analyzing spatial structure in a multi-variable process. The local estimate will change every time \mathbf{X} is changed, which fixes our earlier concern with the Wartenberg and Lee estimators that the local statistic relating y and x_1 ignores all other covariates. The auxiliary regression assigns all of the *direct* spatial variation in $\mathbf{W}y$ as arising solely from \mathbf{X} first. Further, the auxiliary approach can be extended to an arbitrary number of variates and outcomes, providing a useful multivariate-conditional analogue to the Wartenburg

⁵Derivation of this statistic is provided in supplemental material.

Name		Statistic	Implied Outlier Comparison
Univariate	\mathbf{I}_y	$y \circ \mathbf{W}y$	Is this house’s rent unusual for its area?
Bivariate	\mathbf{I}_{xy}	$x \circ \mathbf{W}y$	Is this lot size unusual for house rent in this area?
Spatial ρ	\mathbf{L}_{xy}	$\mathbf{W}x \circ \mathbf{W}y$	Are this area’s lot sizes unusual for this area’s rents?
Partial	$\mathbf{I}_{y x}$	$\mathbf{I}_y - \rho_{xy}\mathbf{I}_{xy}$	Is this house’s rent unusual for its area given its lot size?
Auxiliary	$\mathbf{I}_{x \rightarrow y}$	$\mathbf{I}_{y x} - \rho_{xy}\mathbf{I}_{yx} - \rho_{xy}^2\mathbf{I}_x$	Is this house’s premium unusual for premia in this area?

Table 1: Statistics discussed in this paper, alongside an outlier comparison that the statistic implies.

estimator, as Eckardt and Mateu (2021) show. However, a consequence of the partial regression strategy in the case of a *specific variable of interest* (like y) means that the auxiliary regression removes *all* variation attributable to \mathbf{X} , even that which arises indirectly through $\mathbf{W}\mathbf{X}$. As an illustration, we can return to our rent example where an $\mathbf{I}_{x \rightarrow y}$ outlier would represent an unusual house *premium*. If the predicted rent for a given house is \hat{y}_i , then the *premium* for the house—the extra you would have to pay over-and-above the estimated value—is the residual e_i . Because the auxiliary regressive strategy compares e with $\mathbf{W}e$, it measures how the estimated premium compares to *nearby premia*. It is also not clear how to apply this correction in the case of a bivariate analysis relating some x_p variate to $\mathbf{W}y$, because e becomes the “new” outcome of interest. This means that both the direct influence of \mathbf{X} and the *indirect* influence of $\mathbf{W}\mathbf{X}$ are carried through to our local statistic. The auxiliary Moran statistics do not compare a size-adjusted rent to other rents, and comparing \hat{y} to $\mathbf{W}\hat{y}$ ultimately boils back down to a re-scaled version of \mathbf{I}_x . Thus, the auxiliary regressive local Moran statistic allows us to understand the relationship between premia, but it *does not* examine rent itself.

The Partial Local Moran Statistic

By this point, we have discussed four strategies, each with an increasing usefulness in the case of confounded local inference. In particular, the auxiliary regression strategy is most promising, but is still dissatisfying because it shifts the analysis to consider only the part of y that is independent of \mathbf{X} , changing the variable of interest from y to e . As a result, this “brings along” the effect of $\mathbf{W}\mathbf{X}$ into the final local statistic, as seen in Equation 9. In this way, we can also think of the auxiliary regression strategy as a “full conditional” local Moran statistic, since it fully conditions-out the variation due to \mathbf{X} .

Recognizing this, we have a useful counterpoint for a *partial local Moran* statistic that keeps the variable of interest the same and omits indirect corrections. In the following sections, we show that the same partial local Moran statistic can be derived in two different ways. To start, we will show how it can be obtained from a single Moran-form multiple regression, instead of the matrix of pairwise regressions (as in Wartenberg (1985) & Lee (2001) illustrated in Figure 3) or a two-step auxiliary regression method. The local statistics that result from this are outlined in Table 1, alongside the other common Moran-type statistics. After this, we show how these all can be obtained *graphically* using the paths in the network of direct and indirect associations between y and

\mathbf{X} through \mathbf{W} . This provides a useful conceptual tool to reason about the “right” local statistic for any scenario of interest. After these two discussions, we then show the different ways of correcting for confounding in an empirical analysis of rent in a rapidly-gentrifying neighborhood that has been a testbed for rent stabilization and tenant rights policy. In this empirical example, I discuss four properties in detail and describe how they are treated differently by each statistic. I close by discussing when each kind of correction be appropriate, expounding on the outlier comparisons presented in Table 1.

A Mathematical Explanation

To define the partial local Moran statistic, let’s focus on just one z-standardized outcome variate y and confounding variable x that influences y . We can collect these two into the following design matrix:

$$\mathbf{D} = \begin{bmatrix} y & x \end{bmatrix} \quad (10)$$

Using this design matrix, we can state a *Moran-form multiple regression* as:

$$\mathbf{W}y = \mathbf{D}m + \epsilon \quad (11)$$

where m is a typical vector of slopes in a multiple regression. As a consequence, the first component of m corresponds to the relationship between y and $\mathbf{W}y$ holding the remaining \mathbf{X} constant.⁶ Given this, we can use the ordinary least squares estimate of m as a “partial global Moran estimator:”

$$\hat{m} = (\mathbf{D}'\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}y \quad (12)$$

To exchange the inner product between $\mathbf{D}'\mathbf{W}y$ for an elementwise product as we did in the univariate case we must⁷ tile $\mathbf{W}y$ to match the shape of \mathbf{D} . With this, we can state the local partial Moran statistic as follows:

$$\mathbf{M} = (\mathbf{D}'\mathbf{D})^{-1}(\mathbf{D} \circ \mathbf{W}\mathbf{R})(N - 1) \quad (13)$$

$$\mathbf{W}\mathbf{R} = \mathbf{W} \begin{bmatrix} y & \cdots & y \end{bmatrix}$$

The partial local Moran matrix \mathbf{M} contains the local statistics relating $\mathbf{W}y$ to the variables of interest (y, \mathbf{X}) after correcting for their internal co-variation. In this form, \mathbf{X} can include many variables and \mathbf{M} expands accordingly. Each column of \mathbf{M} contains a different vector of local partial cross- or auto-correlation statistics corresponding to different correlation-weighted combinations of univariate or bivariate local Moran statistics. For example, the case with two variates (x and y) results in the

⁶This ignores the simultaneity between y and $\mathbf{W}y$ just like the univariate Moran-form regression.

⁷Swapping $\mathbf{D}'\mathbf{W}y$ for $\mathbf{D} \circ \mathbf{W}y$ would create an $N \times P$ matrix sitting in the midst of a $P \times P$ covariance matrix $([\mathbf{D}'\mathbf{D}]^{-1})$ and an $N \times 1$ spatial lag of y . Tiling removes this shape mismatch.

following partial local Moran statistic⁸ for y given x :

$$\mathbf{I}_{y|x} = \frac{N-1}{N(1-\rho_{xy}^2)} [y \circ \mathbf{W}y - \rho_{xy}x \circ \mathbf{W}y] \quad (14)$$

For most analyses, this first vector of \mathbf{M} is the most useful: it is the conditional local indicator of association between y and $\mathbf{W}y$ after controlling for the *direct* association between x and y . Intriguingly, this is also the first two terms from the auxiliary/full conditional local Moran statistic (Eq. 9); other terms that depend on $\mathbf{W}x$ are dropped from this “partial” estimator. In addition, if y is already independent of x , then the statistic reduces to the standard univariate local Moran statistic.

The cross-variable analogue to Wartenberg (1985) is contained in the remaining columns of \mathbf{M} corresponding to the \mathbf{X} variates. This provides the cross-correlation between each x and $\mathbf{W}y$ accounting for x ’s direct co-variation with y . As might be expected, $\mathbf{I}_{x|y}$ takes a similar form to $\mathbf{I}_{y|x}$, reflecting the same correction strategy applied to the relationship between y and $\mathbf{W}y$:

$$\mathbf{I}_{x|y} = \frac{N-1}{N(1-\rho_{xy}^2)} [x \circ \mathbf{W}y - \rho_{xy}y \circ \mathbf{W}y] \quad (15)$$

When x and y are independent, this returns the standard bivariate local Moran statistic. But, when x and y are related, this corrects the co-variation between x and $\mathbf{W}y$ that arises from the direct relationship between x and y . This has no analogue in the “full” conditional Moran statistic, but provides a bivariate indicator of spatial association that corrects for the direct relationship between x and y .

A Graphical Explanation

These kinds of corrections to the univariate statistics come naturally, too, if we plot out the different kinds of associations between the variables and their lags. I show this in Figure 4. In the first panel of Figure 4, we see the typical univariate statistics used in spatial analysis that relate two variates (x, y) or their spatial lags ($\mathbf{W}x, \mathbf{W}y$) at site 1 with neighboring sites 2 and 3. Each statistic reflects site 1’s local contribution to a global measure of association. Some links, such as the correlations (ρ_{xy}, L_{xy}) are symmetric, meaning that the association between the two factors can be reversed and the same result obtained. However, Moran-type statistics are *not* reversible in the most common case, when row-standardized weights are used. Hence, symmetric associations are represented as arrows in both directions, while asymmetric relations have arrows in only one direction.

To illustrate where each term in the partial (and full) local Moran statistics arises, we start at the outcome of interest $\mathbf{W}y_1$ and trace each combination of paths back to our variable of interest (either y or x). We take each directed path exactly one time. For example, the second panel shows our new Partial Moran statistic $I_{y|x}$. We obtain the terms shown in Equation 14 by tracing each path *backwards* from $\mathbf{W}y_1$ to y_1 . This shows that the partial estimator uses two different pathways

⁸This is derived in supplemental material for a bivariate example and for three variables.




Figure 4: Diagram of relationships between different variables considered in local Moran-type statistics. Colors indicate each “pathway” that the multivariate conditional statistics use. In the partial conditional local Moran (middle two panes), $\mathbf{W}x_1$ is ignored, and all links through it are removed from the graph. And, in the full conditional local Moran (right), the top red ρ_{xy} edge is used by both the orange and the maroon pathways.

through which $\mathbf{W}y_1$ can be influenced by y_1 . The first “direct” path provides I_y , the $y \circ \mathbf{W}y$ term in Equation 14. The second “indirect” pathway through I_{xy} and ρ_{xy} reflects a correction of I_y due to the association between x and y , and x and $\mathbf{W}y$. Together, these are the two terms in $I_{y|x}$. Likewise, the $I_{x|y}$ diagram in the third pane traces the “direct” path I_{xy} and then the “indirect” correction term through y_1 . Notably, all of the links to $\mathbf{W}x_1$ are omitted in the partial local Moran. The “full” auxiliary regression statistic $I_{x \rightarrow y}$, on the right, retains these links. As a consequence, it can be thought of as the $I_{y|x}$ statistic corrected using the two additional indirect paths that flow through $\mathbf{W}x_1$.⁹

Partial and Full Conditional Local Moran Statistics in Practice

This partial local Moran statistic has quite a few desirable properties. First, there is a “ladder” of restriction that occurs when moving from the “full” to “partial” to the existing univariate and bivariate statistics. The “full” statistic listens to all possible direct and indirect pathways between the variates through \mathbf{W} , while the partial statistic ignores any indirect exogenous pathways $\mathbf{W}X$. The univariate/bivariate statistics only consider the direct pathways. As a consequence, each statistic can be recovered from the higher-level statistic with more “open” pathways, and everything reduces back down to the univariate/bivariate spatial auto/cross-correlation statistic when $\rho_{xy} = 0$. Second, these statistics obey the same properties suggested by Anselin (1995): the sum of obtained local statistics is proportional to a clear global measure of association. Third, the global partial Moran statistic is a typical slope from a multiple regression, like the univariate Moran statistic is a slope in a simple regression. This means that each of these statistics is a conditional, not joint, multivariate statistic—inheriting the same *ceteris paribus* understanding of co-variation as in regression

⁹Each pathway in the full $I_{x \rightarrow y}$ statistic is shown separately in the Supplementary Material, for clarity.

generally. Returning to our housing example, a partial local Moran statistic can determine whether a house is unusually expensive in its neighborhood, given its number of rooms. This is distinct from the auxiliary local Moran statistic, which must remove this variation first, and then analyze the spatial structure of what is left over. Fourth, the addition (or removal) of more \mathbf{X} variates will change the parameter covariance matrix ($\mathbf{D}'\mathbf{D}$) and, as a consequence, the local statistics. This resolves the main issue with the existing bivariate local Moran statistics. Fifth, “conditional Moran scatterplots” and cluster/outlier classifications can be recovered from the partial dependence structure: we can regress $\mathbf{W}y$ onto the “corrected” y to obtain the estimated global conditional Moran’s I .¹⁰ While the full conditional statistic changes the outcome of interest, it can also be expressed in this manner. Finally, permutation inference can be used to conduct local statistical tests.¹¹ I demonstrate all of these features below in an empirical example.

The Local Spatial Structure of Rent in Bristol

Bristol is a growing city of around half a million people in the most rapidly-growing region of the United Kingdom. As a consequence of rapid population growth coupled with inconsistent construction of housing, Bristol has experienced unusual and distinctive patterns of rent-driven population displacement (Boddy 2007; Palmer 2019) with clear negative impacts on inhabitants’ lives (Anguelovski et al. 2021). Community concern about the racial, ethnic, and class-dimensions of this displacement have led to repeated attempts by local government to improve the justice and equity of housing outcomes (Rose et al. 2013).¹² One challenge in analyzing the spatial patterns of rent is that housing *stock* varies systematically in most cities, Bristol included. At its most basic, larger properties tend to have higher rents, and may be more common in some areas of the city. The relationship between a property’s rent and the surrounding rent is *confounded by* the size of the properties being rented. In order to identify spatial clusters in rent that indicate (un)affordability, we need to somehow account for the relationship between flat size and rent. To illustrate our new knowledge about confounded local inference, this provides a very useful and clear empirical example.

So, we consider the database of all rental listings made between 2018 and 2021 on Zoopla,¹³ a main rental market data aggregator. This provides us with information at the rental listing level, including the number of bedrooms in the rental, the weekly advertised rent for the property, a large amount of text about the property, indicative images, and the energy rating of the property. The Bristol City Council and community stakeholders seek to understand the impact of “living rent” policies, and is thus interested in understanding what areas may be most affected by these policies. Thus, understanding the spatial structure of the rental market requires us to *control for* some of the

¹⁰A full treatment of this is provided in the supplemental material.

¹¹Although, closed-form estimators may come available following the strategies outlined in Sauer et al. (2021). The exact procedure for local permutation testing on conditional Moran statistics is shown in supplement. Implementations of the estimators will be made available in production-ready R (Bivand 2022) and Python (Rey et al. 2022) packages after submission. For peer review, estimators are provided alongside the submission.

¹²This includes the government commission exploring “living rent” policies that supports this work.

¹³This data is provided upon request from the Urban Big Data Center.

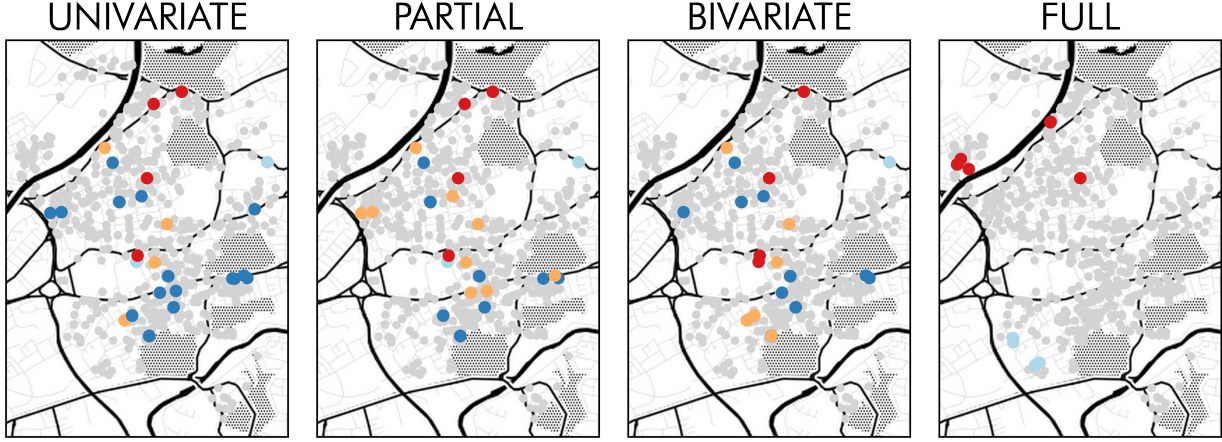


Figure 5: Map of all statistically-significant local statistics ($p \leq 0.05$) for rental properties in East Bristol, using similar coloring as to Figure 1.

potential pricing inputs to a flat’s advertised rent in order to analyze its local spatial structure.

Therefore, in the following example, we present a simplified analysis focusing on a single extremely influential determinant of rent: property size. Let the weekly rent in pounds be y , and the number of bedrooms listed with the property be x . The correlation between advertised rent and number of bedrooms is 0.84, suggesting a very strong relationship. Larger properties are often split by multiple tenants who pool their payments or simply reflect a tenant’s preference (or need) to rent “more” property. This relationship itself is well-understood in the analysis of affordability, but it be challenging to account for its spatial variation *at the same time* as the variation in rent. We can see this variation in Figure 2). This variation can be thought of econometrically in terms of *choice sets* (Bruch and Swait 2019), where individuals seeking to locate in a particular area consider alternative rental properties with similar traits that are nearby. Thus, price competition occurs in geographic *clusters* with similar attributes (Wolf 2021). To model these local alternatives, we use a Delaunay triangulation between the property locations.¹⁴

The statistically significant local Moran clusters and outliers obtained by conditional permutation inference (e.g. Sauer et al. 2021) are also colored in the scatterplots and mapped in Figure 5. The classifications are also reproduced in Table 2, which shows how the classifications that *consider* property size differ from the univariate analysis considering only rent. Between these two exhibits, we can see that most observations have the same classification across the univariate, bivariate, and partial local Moran statistics, although quite a few “coldspots” move to “high outliers” in the partial statistic. For the fully-conditional statistic, only *one* univariate hotspot remains a fully-conditional hotspot. This demonstrates how this correction may be over-eager, especially when ρ_{xy} is large. In contrast, the *partial* correction still results in a few univariate coldspots falling from sig-

¹⁴The substantive results we show are consistent with sparse k-nearest neighbor graphs as well ($k \in \{1, 2, \dots, 10\}$). Kernel functions were also explored and exhibit the same qualitative behavior, but tended to over-smooth the W_{ij} values, so the Delaunay triangulation was chosen to make the examples clear.

Univariate	N.S.	Hotspot	Low Outlier	Coldspot	High Outlier	Comparator
Not Significant	1412	0	0	0	0	Bivariate
Hotspot	2	4	0	0	0	
Low Outlier	1	1	2	0	0	
Coldspot	7	0	0	21	4	
High Outlier	0	0	0	0	4	
Not Significant	1412	0	0	0	0	Partial
Hotspot	0	6	0	0	0	
Low Outlier	0	0	4	0	0	
Coldspot	4	0	0	20	8	
High Outlier	1	0	0	0	3	
Not Significant	1402	6	0	4	0	Full
Hotspot	5	1	0	0	0	
Low Outlier	4	0	0	0	0	
Coldspot	32	0	0	0	0	
High Outlier	4	0	0	0	0	

Table 2: Crosstabulation of local Moran classifications for listings shown in Figure 5. “N.S.” stands for a local statistic that is not significant at the 95% level.

nificance, but also shows a few observations that shift from coldspots (low rent observations in low rent areas) over to high outliers. This suggests that the partial statistic is better able to retain the information about rent directly while also correcting for property size.

The plots in Figure 6 illustrate this further. There, I show how four example properties change between the different kinds of local Moran statistics. The first annotated example property is the “lux 8-bed” property marked by the gold star. It is a large newly-renovated property that was re-developed to attract a shared tenancy of young professionals. The “pricey studio,” marked by the red “x,” is a very small but expensive single-occupant property. The “2-bed steal” reflects a two-bedroom property that was advertised as a property with relatively low rent given other rents in the area for similar properties. The “ideal family home” was similarly advertised as a relatively cheap property in area for a family to rent. The rental properties are also located in very different parts of the neighborhood, so each example property has a distinct local choice set while still reflecting samples from the same property market area.

Each of these properties occupies a different position in the distribution of flat sizes, as can be seen in the “Pearson” facet of Figure 6. The Luxe 8-Bed and Pricey Studio are unusually expensive for their property sizes, since they occur well above the overall trend line for the rest of the dataset. But, in the univariate local Moran scatterplot, the studio appears unremarkable, suggesting that its rent is not unusually high considering the rents in its area. In the partial and full conditional Moran scatterplots, the Pricey Studio emerges as a uniquely expensive property in a uniquely expensive area. This again is expected, because studios are cheaper on average because they represent “less

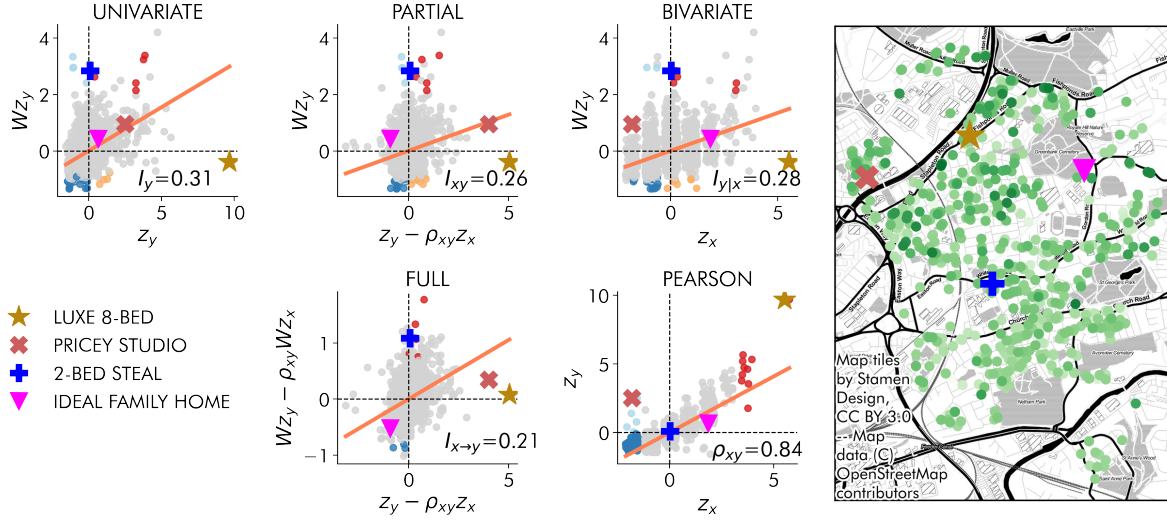


Figure 6: Example properties selected in East Bristol, chosen to demonstrate how the conditional Moran statistics change depending on the local relationship between rent and flat size.

property” being rented. The Luxe 8-Bed is classified as a “hotspot” in the Full Conditional Moran because of the relatively high rent premia in the area, whereas the generally lower *rent* in this area makes it an outlier in the other three local Moran classifications. The Pricey Studio, however, becomes a high rent cluster in both conditional statistics. Further, all of the Moran Scatterplots demonstrate that the 2-Bed Steal has fairly typical rent but is in a very expensive area. Because it is surrounded by many other two bedroom properties, it keeps its relative position fairly constant across all of the Moran Scatterplots, conditional, univariate, or bivariate.

Finally, the Ideal Family Home shifts around substantially, and provides a useful point of comparison for the 2-Bed Steal. This listing has relatively low rent for four bedroom properties but is much further east, putting it substantially further away from transit links and the city center. As a consequence, rent premia are unusually low in this area. But, relatively-larger rental properties are *also* in this area. Therefore, rents in this area are close to the map average, even though the properties are larger. This means that the property shifts from being in the “hotspot” quadrant to being a “low-high” outlier for the partial conditional Moran since its size-adjusted rent is low relative to surrounding unadjusted rents. This might mean the property is unusually large or unusually cheap for its area. In contrast, becomes “coldspot” in the full conditional Moran statistic because it has a low premium and is located in a low-premium area. This adjusts the classification again, because the other properties in the area are also cheap for their size. This shows clearly that the different corrections clearly result in different interpretations, as the surrounding flat sizes do not enter into the model for rent in the univariate or partial statistic, but do in the full statistic.

Discussion & Conclusion

Local spatial analysis is a critical part of exploratory spatial data analysis. Chief among exploratory spatial data analysis methods are the local indicators of spatial association (Anselin 1995). One of the most commonly-used methods in this family is the univariate local Moran statistic, which we extend with new conditional estimators that can account for *confounded* local inference. In particular, the partial and full conditional local Moran statistics provide useful new methods for analyzing the local structure of data when there is more than one variable being studied. They extend the univariate statistic, relating the spatial distribution of the original attribute under study to a covariance-weighted composite of y and potentially-confounding X . The full conditional statistic also included the indirect influence of WX . I develop a method for inference for these local statistics that relies on similar conditional map randomization strategies. Then, I show how the analysis plays out for renting in an area of east Bristol, discussing case studies that change in interpretation from statistic to statistic. Critically, we have shown how these estimators arise from the network of associations between these different factors, and hope this kind of network-based thinking is useful for future scholars conducting local analysis.

One clear question that may arise from this is: *what is the right statistic to use?* Hopefully, we have already indicated few clear arguments that can help answer this question. We agree with Lee (2001) that the original bivariate Moran estimator models an unusual association that is unlikely to be of direct interest in most application. Instead, the partial estimator should nearly always be used in this case, even over the Lee statistic, because it more completely corrects for the direct and indirect associations between y and X , and can extend beyond bivariate relationships. For choosing between the full and partial conditional local Moran statistics, the partial statistic is most applicable when the outcome of interest needs to remain the same. If the analysis remains useful when focused on the residual e instead of y , then the full statistic can be used. The rent example demonstrates this: the partial statistic is most useful in this case, since it remains a comparison of the size-adjusted rent to surrounding rents. However, the full correction is still useful to understand whether the premium a renter pays is unusual. Any disagreement between the partial and full statistics arises solely from whether we consider the premium or the rent as the outcome variate.

Further, as regression-derived statistics, typical regression *assumptions* will matter for *all* of these statistics considered. Heteroskedasticity and/or nonlinearity may result in odd or unusual local estimates, regardless of any correction being applied. In fact, the relationship analyzed in this paper is slightly heteroskedastic, like most data on prices. The influence of distributional corrections on local statistics remains unknown, though, in general. And, while we treated the number of bedrooms as a continuous variable, these approaches could also represent this using a “dummy” variable to model group-wise heterogeneity. Thus, using a local statistic should be treated with *as much care as* using a full multivariate regression model, and the same considerations should, at minimum, be made.

In general, both of the conditional Moran approaches inherit from the same Moran-form regres-

sion. The estimators are all intrinsically linked: the univariate statistic is recovered from the partial estimator, and the partial estimator from the full conditional estimator. They retain Moran Scatterplots as an exploratory visualization technique, empowering the existing exploratory spatial data analysis visualization methods such as linking and brushing of map and scatterplot views (Anselin and Rey 2014) and quadrant classifications that are popular in conceptualization local spatial relationships. Thus, the partial and the full conditional local Moran statistic can help make sense of our multivariate world.

References

- Anguelovski, Isabelle, Helen V. S. Cole, Ella O'Neill, Francesc Baró, Panagiota Kotsila, Filka Sekulova, Carmen Pérez del Pulgar, Galia Shokry, Melissa García-Lamarca, Lucia Argüelles, James JT Connolly, Jordi Honey-Rosés, Antonio López-Gay, Mario Fontán-Vela, Austin Matheney, Emilia Oscilowicz, Andrew Binet, and Margarita Triguero-Mas (2021). "Gentrification Pathways and Their Health Impacts on Historically Marginalized Residents in Europe and North America: Global Qualitative Evidence from 14 Cities". *Health & Place* 72, p. 102698. ISSN: 1353-8292. DOI: 10.1016/j.healthplace.2021.102698. (Visited on 10/05/2022).
- Anselin, Luc (1995). "Local Indicators of Spatial Association-LISA". *Geographical Analysis* 27.2, pp. 93–115.
- (1996). "The Moran Scatterplot as an Exploratory Spatial Data Analysis Tool to Assess Local Instability in Spatial Association". *Environmental Modeling with GIS*. Ed. by M. M. Fischer, H. J. Scholten, and D. Unwin. Oxford: Oxford University Press, pp. 454–469.
- Anselin, Luc and Daniel A Griffith (1988). "Do Spatial Effects Really Matter in Regression Analysis?" *Papers Regional Science* 65, pp. 11–34.
- Anselin, Luc and Xun Li (2020). "Tobler's Law in a Multivariate World". *Geographical Analysis* 52, pp. 494–510. ISSN: 1538-4632. DOI: 10.1111/gean.12237. (Visited on 05/13/2020).
- Anselin, Luc and Sergio J. Rey (2014). *Modern Spatial Econometrics in Practice, a Guide to GeoDa, GeoDaSpace, and PySAL*. Chicago, IL: GeoDa Press.
- Anselin, Luc, Ibnu Syabri, and Oleg Smirnov (2002). "Visualizing Multivariate Spatial Correlation with Dynamically Linked Windows". *Proceedings*. Santa Barbara, CA.
- Bivand, Roger (2022). "R Packages for Analyzing Spatial Data: A Comparative Case Study with Areal Data". *Geographical Analysis* 54.3, pp. 488–518. ISSN: 0016-7363, 1538-4632. DOI: 10.1111/gean.12319. (Visited on 10/18/2022).
- Boddy, Martin (2007). "Designer Neighbourhoods: New-Build Residential Development in Non-metropolitan UK Cities—The Case of Bristol". 39, pp. 86–105.
- Bruch, Elizabeth and Joffre Swait (2019). "Choice Set Formation in Residential Mobility and Its Implications for Segregation Dynamics". *Demography* 56.5, pp. 1665–1692. ISSN: 0070-3370, 1533-7790. DOI: 10.1007/s13524-019-00810-5. (Visited on 03/07/2022).

- Dray, Stéphane and Thibaut Jombart (2011). "Revisiting Guerry's Data: Introducing Spatial Constraints in Multivariate Analysis". *The Annals of Applied Statistics*, pp. 2278–2299.
- Eckardt, Matthias and Jorge Mateu (2021). "Partial and Semi-Partial Statistics of Spatial Associations for Multivariate Areal Data". *Geographical Analysis* 53.4, pp. 818–835. ISSN: 0016-7363, 1538-4632. DOI: 10.1111/gean.12266. (Visited on 10/26/2023).
- Filzmoser, Peter, Anne Ruiz-Gazen, and Christine Thomas-Agnan (2014). "Identification of Local Multivariate Outliers". *Statistical Papers* 55.1, pp. 29–47.
- Fotheringham, A Stewart (1993). "On the Future of Spatial Analysis: The Role of GIS". *Environment and Planning A* 25, pp. 30–34.
- Fotheringham, A Stewart and Chris Brunsdon (1999). "Local Forms of Spatial Analysis". *Geographical Analysis* 31.4, pp. 340–358.
- Getis, Arthur and J. Keith Ord (1992). "The Analysis of Spatial Association by Use of Distance Statistics". *Geographical Analysis* 24.3, pp. 189–206.
- (1996). "Local Spatial Statistics: An Overview". *Spatial Analysis: Modelling in a GIS Environment* 374.
- Harris, Paul, Chris Brunsdon, Martin Charlton, Steve Juggins, and Annemarie Clarke (2014). "Multivariate Spatial Outlier Detection Using Robust Geographically Weighted Methods". *Mathematical Geosciences* 46.1, pp. 1–31.
- Jones, Kelvyn (1991). "Specifying and Estimating Multilevel Models for Geographic Research". *Transactions of the Institute of British Geographers* 16, pp. 148–159.
- Lee, Sang-Il (2001). "Developing a Bivariate Spatial Association Measure: An Integration of Pearson's r and Moran's I ". *Journal of Geographical Systems* 3.4, pp. 369–385.
- LeSage, James P and R. Kelley Pace (2009). *Introduction to Spatial Econometrics*. Boca Raton, FL: CRC Press.
- (2014). "The Biggest Myth in Spatial Econometrics". *Econometrics* 2.4, pp. 217–249. ISSN: 2225-1146. DOI: 10.3390/econometrics2040217. (Visited on 05/16/2022).
- Palmer, Henry (2019). *Voices of Bristol: Gentrification and Us*. Arkbound. ISBN: 978-1-912092-90-1.
- Rey, Sergio J., Luc Anselin, Pedro Amaral, Dani Arribas-Bel, Renan Xavier Cortes, James David Gaboardi, Wei Kang, Elijah Knaap, Ziqi Li, Stefanie Lumnitz, Taylor M. Oshan, Hu Shao, and Levi John Wolf (2022). "The PySAL Ecosystem: Philosophy and Implementation". *Geographical Analysis* 54.3, pp. 467–487. ISSN: 1538-4632. DOI: 10.1111/gean.12276. (Visited on 08/18/2022).
- Rogerson, Peter A. (2022). "Scan Statistics Adjusted for Global Spatial Autocorrelation". *Geographical Analysis* 54.4, pp. 739–751. ISSN: 0016-7363, 1538-4632. DOI: 10.1111/gean.12301. (Visited on 10/13/2022).
- Rose, Damaris, Annick Germain, Marie-Hélène Bacqué, Gary Bridge, Yankel Fijalkow, and Tom Slater (2013). "'Social Mix' and Neighbourhood Revitalization in a Transatlantic Perspective: Comparing Local Policy Discourses and Expectations in Paris (France), Bristol (UK) and Montréal (Canada)". *International Journal of Urban and Regional Research* 37.2, pp. 430–450. ISSN: 1468-2427. DOI: 10.1111/j.1468-2427.2012.01127.x. (Visited on 10/05/2022).

- Sauer, Jeffery, Taylor M. Oshan, Sergio J. Rey, and Levi John Wolf (2021). "The Importance of Null Hypotheses: Understanding Differences in Local Moran's under Heteroskedasticity". *Geographical Analysis*. DOI: 10.1111/gean.12304.
- Wartenberg, Daniel (1985). "Multivariate Spatial Correlation: A Method for Exploratory Geographical Analysis". *Geographical Analysis* 17.4, pp. 263–283. (Visited on 02/08/2016).
- Wolf, Levi John (2021). "Spatially-Encouraged Spectral Clustering: A Technique for Blending Map Typologies and Regionalization". *International Journal of Geographical Information Science* 35.11, pp. 2356–2373. ISSN: 1365-8816, 1362-3087. DOI: 10.1080/13658816.2021.1934475. (Visited on 10/14/2022).
- Wolf, Levi John, Taylor M. Oshan, and A Stewart Fotheringham (2018). "Single and Multiscale Models of Process Spatial Heterogeneity". *Geographical Analysis* 50.3, pp. 223–246. ISSN: 1538-4632. DOI: 10.1111/gean.12147. (Visited on 10/18/2018).

Supplemental Material

Conditional Moran: Two-dimensional example

Since the analytical expression for a matrix inverse with more than two rows gets verbose quickly, we proceed using a single additional covariate, x , to keep the presentation manageable. From the full statement in Equation 13, we replace \mathbf{D} and \mathbf{R} with their component vectors:

$$\mathbf{M} = \begin{bmatrix} y & x \end{bmatrix} \circ \mathbf{W} \begin{bmatrix} y & y \end{bmatrix} \begin{bmatrix} y'y & y'x \\ x'y & x'x \end{bmatrix}^{-1} (N-1) \quad (16)$$

Noting that y and x are unit-variance and centered, $y'x = x'y = n\rho_{xy}$, the typical aspatial correlation between y and x , and $x'x$ and $y'y$ are both n . Therefore, we can factor this out to obtain a correlation matrix:

$$\mathbf{M} = \begin{bmatrix} y & x \end{bmatrix} \circ \mathbf{W} \begin{bmatrix} y & y \end{bmatrix} \begin{bmatrix} 1 & \rho_{xy} \\ \rho_{xy} & 1 \end{bmatrix}^{-1} \frac{N-1}{N} \quad (17)$$

Now that we have obtained this expression, we can restate the inverse of $(\mathbf{D}'\mathbf{D})^{-1}$ in terms of correlations:

$$\mathbf{M} = \begin{bmatrix} y & x \end{bmatrix} \circ \mathbf{W} \begin{bmatrix} y & y \end{bmatrix} \begin{bmatrix} 1 & -\rho_{xy} \\ -\rho_{xy} & 1 \end{bmatrix} \frac{N-1}{N(1-\rho_{xy}^2)} \quad (18)$$

Thus, the local structure of y and its surroundings, jointly determined by the relationship between x and y , is:

$$\mathbf{M}_1 = \mathbf{I}_{y|x} = \begin{bmatrix} y \circ \mathbf{W}y - \rho_{xy}x \circ \mathbf{W}y \end{bmatrix} \frac{(N-1)}{N(1-\rho_{xy}^2)} \quad (19)$$

This is the conditional Moran for y , which embeds the relationship between y and itself in space, estimated jointly with another variate, x . Practically speaking, the first term reflects all of the information about the spatial and aspatial correlation between terms. The first part of this first term, $y \circ \mathbf{W}y$, is proportional to the univariate local Moran's I_i . The second term is proportional to the Wartenberg (1985) cross-association measure ($x \circ \mathbf{W}y$) rescaled by the aspatial association between x and y . In this sense, the first component is the “direct” spatial relationship capturing the component originally modeled in the univariate Moran, and the second corrects for the “indirect” relationships between x , y , and $\mathbf{W}y$. If x and y were independent, we would simply have the usual univariate local Moran statistic.

There is also an analogue to the bivariate Moran's I relating x and $\mathbf{W}y$, but one which also accounts for the fact that y is also related to $\mathbf{W}y$. This is the second column of \mathbf{M} , relating $\mathbf{W}y$ to x . Its local terms have the form:

$$\mathbf{M}_2 = \mathbf{I}_{x|y} = \begin{bmatrix} x \circ \mathbf{W}y - \rho_{xy}y \circ \mathbf{W}y \end{bmatrix} \frac{(N-1)}{N(1-\rho_{xy}^2)} \quad (20)$$

If y were independent of x in this case, we would recover the standard bivariate local Moran statis-

tic. But, since this is rarely the case, we need to correct for the association between y and x directly. This again provides measures the association between x and $\mathbf{W}y$, estimated alongside the relationship between x and y as well as that between y and $\mathbf{W}y$. The simultaneity resolves our conceptual issue with existing bivariate local Moran statistics: the relationship between x and $\mathbf{W}y$ is estimated while acknowledging that y is observed along with x and $\mathbf{W}y$.

Conditional Moran: More Than Two Dimensions

This derivation can also be arbitrarily extended to more than one y and one x , meaning that an arbitrary number of additional “spatial confounders” may be included in \mathbf{D} . However, this presentation gets verbose quickly due to the symbolic expressions for the matrix inverse. Let an additional covariate, z , be considered. This makes the design matrix:

$$\mathbf{D} = \begin{bmatrix} y & x & z \end{bmatrix} \quad (21)$$

The same general estimator from Equation 13 applies:

$$\mathbf{M} = (\mathbf{D} \circ \mathbf{WR})(\mathbf{D}'\mathbf{D})^{-1} \quad (22)$$

But now, the equation for $(\mathbf{D}'\mathbf{D})^{-1}$ becomes more complex. We can still define this in terms of the correlation matrix:

$$(\mathbf{D}'\mathbf{D}) = \begin{bmatrix} 1 & \rho_{yx} & \rho_{yz} \\ \rho_{xy} & 1 & \rho_{xz} \\ \rho_{zy} & \rho_{zx} & 1 \end{bmatrix} \frac{1}{N} \quad (23)$$

However, to take a symbolic inverse of this matrix, we must first recognize the relationship between a matrix inverse, its determinant, and *adjugate* matrices of $\mathbf{D}'\mathbf{D}$:

$$(\mathbf{D}'\mathbf{D})^{-1} = \text{adj}(\mathbf{D}'\mathbf{D}) \frac{1}{|\mathbf{D}'\mathbf{D}|} = \begin{bmatrix} C_{yy} & -C_{yx} & C_{yz} \\ -C_{xy} & C_{xx} & -C_{xz} \\ C_{zy} & -C_{zx} & C_{zz} \end{bmatrix}' \frac{1}{|\mathbf{D}'\mathbf{D}|} \quad (24)$$

Each “cofactor” $C_{.,.}$ is a determinant of a sub-matrix of $\mathbf{D}'\mathbf{D}$. For example, $C_{y,z}$ is the determinant of the matrix left over when removing the row corresponding to y and the column corresponding to z in $\mathbf{D}'\mathbf{D}$ —in this case dropping the first row and last column of $\mathbf{D}'\mathbf{D}$:

$$C_{y,x} = \begin{vmatrix} \rho_{xy} & 1 \\ \rho_{zy} & \rho_{zx} \end{vmatrix}$$

When collected together, these are often called the *cofactor* matrix. Since $\mathbf{D}'\mathbf{D}$ is symmetric, the cofactor matrix is as well. And, because the diagonal terms all have the same form, we just need

four terms to define this matrix inverse symbolically:

$$C_{yy} = 1 - \rho_{xz}^2 \quad (25)$$

$$-C_{yx} = \rho_{yz}\rho_{xz} - \rho_{xy} \quad (26)$$

$$C_{yz} = \rho_{xy}\rho_{xz} - \rho_{yz} \quad (27)$$

$$-C_{xz} = \rho_{xy}\rho_{zy} - \rho_{xz} \quad (28)$$

Since the correlations are undirected, the ordering of the subscripts is arbitrary. So, restating the matrix inverse symbolically in terms of correlations and the determinant, the full equation becomes:

$$\mathbf{M} = \begin{bmatrix} y & x & z \end{bmatrix} \circ \mathbf{W} \begin{bmatrix} y & y & y \end{bmatrix} \begin{bmatrix} 1 - \rho_{xz}^2 & \rho_{yz}\rho_{xz} - \rho_{xy} & \rho_{xy}\rho_{xz} - \rho_{yz} \\ \rho_{yz}\rho_{xz} - \rho_{xy} & 1 - \rho_{yz}^2 & \rho_{xy}\rho_{zy} - \rho_{xz} \\ \rho_{xz}\rho_{xy} - \rho_{yz} & \rho_{xy}\rho_{xz} - \rho_{xz} & 1 - \rho_{xy}^2 \end{bmatrix} \frac{N-1}{N|\mathbf{D}'\mathbf{D}|} \quad (29)$$

Focusing only on the first column of \mathbf{M} that corresponds to y , we obtain two kinds of terms for components relating the three covariates to $\mathbf{W}y$. Ignoring the scaling factor, this is $\mathbf{I}_{y|xz}$:

$$\mathbf{M}_1 = \mathbf{I}_{y|xz} \propto \left[(1 - \rho_{xz}^2)(y \circ \mathbf{W}y) + (\rho_{yz}\rho_{xz} - \rho_{xy})(x \circ \mathbf{W}y) + (\rho_{xz}\rho_{xy} - \rho_{yz})(z \circ \mathbf{W}y) \right] \quad (30)$$

Breaking this down, the first term refers to the re-scaled relationship between y and $\mathbf{W}y$. The remaining terms reflect both first-order relationships, such as ρ_{xy} , that relate a variable directly to y , and second-order relationships, such as $\rho_{yz}\rho_{xz}$, that reflect a “mediated” relationship between y and z through x .

The “indirect” terms reflect the second-order correlations, such as $\rho_{xy}\rho_{xz}$, with x or z . In contrast, the direct terms involve x or z scaled solely by $\rho_{.,y}$ or involve $y \circ \mathbf{W}y$ itself. As the number of covariates increases, the indirect terms become much more verbose, since the inverse correlation matrix becomes symbolically complicated. But, the fundamental concept of “correcting” the relationship between y & $\mathbf{W}y$ using the direct correlation between additional x & y as well as the indirect correlations will hold at any size of \mathbf{D} .

Classifications from Partial Local Moran Statistics

The most useful aspect of existing local Moran statistics is the classification into areas of local concordance, where the relationship between y and its surroundings ($\mathbf{W}y$) is more positive than expected from the rest of the map, or local discordance, where the relationship between y and its surroundings is more negative than expected from the rest of the map. When the relationship is more concordant than would be expected from the map pattern at large, these areas are identified as spatial “clusters;” likewise, when they are significantly discordant, areas are identified as spatial “outliers.” Thus, all clusters have positive local Moran statistics, and all outliers have negative local Moran statistics, regardless of the significance level used in permutation inference. As shown

above in Figures 1 & 3, this also corresponds to the relationship between y_i , $\mathbf{W}y_i$, and the means of y and $\mathbf{W}y$ in the univariate case.¹⁵

However, since the local conditional Moran statistic involves more than just y and $\mathbf{W}y$, using the same style of classification—assigning observations to cluster/outlier based on y and $\mathbf{W}y$ alone—is problematic. Fortunately, we can construct an appropriate classification by recognizing that the eventual sign of the i th local conditional Moran statistic, $[\mathbf{I}_{y|xz}]_i$, results from the product of the i th spatial lag, $[\mathbf{W}y]_i$ with the i th “corrected” outcome, \tilde{y}_i . The “corrected” term can be obtained by factoring out $\mathbf{W}y$ from Equation 19:

$$\mathbf{I}_{y|x,z} = \left[(1 - \rho_{xz}^2)y + (\rho_{yz}\rho_{xz} - \rho_{xy})x + (\rho_{xz}\rho_{xy} - \rho_{yz})z \right] \circ \mathbf{W}y = \tilde{y} \circ \mathbf{W}y \quad (31)$$

This re-expression makes clear that the factor in brackets on the left, $\tilde{y} = (1 - \rho_{xz}^2)y + (\rho_{yz}\rho_{xz} - \rho_{xy})x + (\rho_{xz}\rho_{xy} - \rho_{yz})z$, is a version of y “corrected” by the joint covariance of y, x , and z .

The classification based on the signs of \tilde{y} and $\mathbf{W}y$ reflect quadrants in the scatterplot of $\mathbf{W}y$ onto \tilde{y} .¹⁶ Going from univariate to the “multivariate” scatters, only lateral movement is possible since $\mathbf{W}y$ is the same in both the single- and multi-variable cases. This means that observations can never transition from being one type of cluster to another, nor from one type of outlier to another; additional data can only move observations from cluster to outlier or outlier to cluster while remaining at the same vertical level in the scatterplot. This is due to the fact that the response variate, $\mathbf{W}y$, remains the same between the two cases.

Partial Local Moran Statistics are different from Auxiliary Regression

To demonstrate the differences between the partial and conditional Moran strategies, let us first simply define the new auxiliary variable e , the residuals from the regression of y on a single confounding variate, x :

$$e = y - \hat{y} = y - (\hat{\beta}_0 + \hat{\beta}_1 x) \quad (32)$$

Because both y and x are mean and variance-standardized, $\beta_1 = \rho_{xy}$ and $\beta_0 = 0$. This is fortunate, because this allows us to directly see the difference between the partial regression strategy and the

¹⁵It is also important to note that the univariate local I_i is not a statement about the relationship between $\widehat{\mathbf{W}y}$ and $\mathbf{W}y$ *per se*. Significant local statistics can be found both above and below the Moran-form regression line in any quadrant, depending the global trend and target y_i . Positive dependence is ubiquitous empirically, so outliers almost always have $\widehat{\mathbf{W}y}_i$ of opposite sign from $\mathbf{W}y_i$. But, for clusters under positive dependence, $\widehat{\mathbf{W}y}_i$ may be above- or below-trend, depending on how extreme y_i is. Many of these can be seen in Figures 1 & 3, where some identified clusters fall on the “wrong side” of the regression line. While most clusters have more extreme observed $\mathbf{W}z_{y_i}$, those with extreme z_{y_i} may not. Thus, classifications based directly on the sign of the residuals for $\mathbf{W}y$ alone are not strictly correct.

¹⁶Visualizing a line on this scatterplot is possible, but the required line is not the partial regression of $\mathbf{W}y$ on \tilde{y} . The correct line is the corresponding global conditional Moran, the slope of the multiple regression, interpreted with the understanding that only changes in y embody changes in \tilde{y} .




Figure 7: Diagram of the four pathways covered by the Auxiliary Moran Regression statistic.

conditional strategy:

$$\hat{\mathbf{I}}_{x \rightarrow y} = (e \circ \mathbf{W}e)(e'e)^{-1}(N-1) \quad (33)$$

$$= (y - \rho_{xy}x) \circ \mathbf{W}(y - \rho_{xy}x) [(y - \rho_{xy}x)'(y - \rho_{xy}x)]^{-1} (N-1) \quad (34)$$

$$\hat{\mathbf{I}}_{x \rightarrow y} = \left[y \circ \mathbf{W}y - \rho_{xy}x \circ \mathbf{W}y - \rho_{xy}y \circ \mathbf{W}x + \rho_{xy}^2 x \circ \mathbf{W}x \right] \frac{(N-1)}{N(1 - \rho_{xy}^2)} \quad (35)$$

We can see that the two are not equal, although they share some similar structures. Specifically, the conditional Moran strategy omits indirect correction due to $\mathbf{W}x$ terms. This may seem surprising, since $\mathbf{W}x$ never enters into the partial regression strategy explicitly. However, $\mathbf{W}x$ leaks into the model because the partial Moran strategy actually analyzes $\mathbf{W}e = \mathbf{W}(y - \rho_{xy}x)$, not $\mathbf{W}y$ directly. In terms of a conceptual causal diagram, the partial Moran statistic “closes” the pathway between $\mathbf{W}x$ and $\mathbf{W}y$, whereas the auxiliary Moran statistic leaves this pathway open. For multivariate \mathbf{X} , the partial regression opens all $\mathbf{W}X$ pathways, too, as shown in This is also clear when considering the four pathways that compose the partial statistic, shown in Figure 7. Only pathway 1 and 2 are included in the partial statistic, while the full statistic includes pathway 3 and 4.

Conditional randomization for conditional Moran statistics.

To conduct inference on the auxiliary Moran statistics, a commonly-used univariate conditional randomization strategy can be adapted: conditional randomization of e is sufficient to do significance testing of $I_{x \rightarrow y}$. In addition, one can also randomize the set of variates for each site, (y_i, X_i) , so long as x_i and y_i are kept together. For the partial Moran statistic, the i th local statistic is only affected by shuffling through the $\mathbf{W}y$ term, so it is sufficient to apply the same univariate strategy in this case to re-sample the $\mathbf{W}y$ component. For example, the bivariate form of the estimator stated in Eq. 19 shows that $N, x'x$ and $|\mathbf{D}'\mathbf{D}|$ are not affected by the shuffling of y and x so long as y_i remains with X_i . Shuffling changes $[\mathbf{W}y^*]_i$, though, since \mathbf{W} is not reshuffled to be synchronized with y or x . This means that the average value of the neighbors near site i will change, even though

the values y_i and x_i do not. Therefore, if y and x are shuffled in sync keeping $y_i = y_i^*$ and $x_i = x_i^*$, then the only term in Equation 19 that changes is $\mathbf{W}y$.