Department for Applied Statistics
Johannes Kepler University Linz

IFAS Research Paper Series
2008-33

Power Calculations for Global and Local Moran's $\mathcal{I}$

Roger Bivand[a], Werner G. Müller[b]
and Markus Reder[c]

February 2008

[a]Department of Economics, Norwegian School of Economics and Business Administration, Bergen, Norway
[b]Department of Applied Statistics, Johannes Kepler University, Linz, Austria
[c]Department of Statistics and Mathematics, University of Economics, Vienna, Austria

Abstract

As in any statistical test, a power analysis can help in assessing the outcomes of whether we are facing global or local spatial dependencies. Although Tiefelsdorf (2000) addresses this point briefly with respect to global Moran's I, it is surprising that practitioners do not seem to have followed this up more widely. One reason may be that the most commonly used spatial analysis and GIS software packages do not support power analysis. Thus, apart from using the code for saddle-point approximation provided in connection with Tiefelsdorf (2002), users have been practically restricted to employing normal approximations.

In this paper, we present an implementation of the exact distributions for global and local Moran's I, following the developments in Tiefelsdorf and Boots (1995), which are integrated into the R-package **spdep** (Bivand et al., 2008). Furthermore, assuming a simultaneous autoregressive spatial data generating scheme, we provide some substantial cases, demonstrating the drawbacks and potential flaws of using the normal approximation in power calculations. Our results confirm the intuitions expressed in Tiefelsdorf (2000), that particularly for local Moran's I, due to the smallness of sets of neighborhoods, this practice may potentially lead to errors of inference. We present an example concerned with Upper-Austrian migration, where using the exact distribution leads to different conclusions.

Keywords: Normal approximation, exact distribution, hot spots, clusters, spatial dependence.

1 Introduction

Modern spatial analysis requires a local viewpoint to complement global analysis (cf. Lloyd (2007)). This is an essential prerequisite for being able to separate influences due to model misspecification from genuinely spatially determined patterns. We require methods of investigation that allow us to try to identify locally particular behavior against the background of global data generation mechanisms. Global model misspecification should remain the prime suspect, but very often a careful set of local analyses can be of great help in showing ways that the global model can be improved.

One measure that has been applied to detect local spatial autocorrelation is local Moran's $\mathcal{I}_i$, developed from global Moran's $\mathcal{I}$. The specific difficulties that are found in using such local indicators of spatial association are partly to do with adjusting probability values to take account of the repeated use of the same data in inference, and secondly to do with inference itself. We will be addressing the inference question, to establish whether or not we can rely on the normal approximation for local Moran's $\mathcal{I}_i$, given the very small neighbour lattice used for computing each statistic.

1.1 Global and Local Moran's $\mathcal{I}$

The Moran's $\mathcal{I}$ statistic is commonly used as an indicator of spatial dependence in a data set. It is defined as a ratio of quadratic forms in the normally distributed regression residuals $\hat{\boldsymbol{\epsilon}} = (I - X(X^TX)^{-1}X^T)y = My$ from a regression on y on X, and

thus of the same structure as the Durbin-Watson statistic for serial autocorrelation, i.e.

$$\mathcal{I}_0 = \frac{\hat{\boldsymbol{\epsilon}}^T \cdot \frac{1}{2} \cdot \left(\boldsymbol{V} + \boldsymbol{V}^T\right) \cdot \hat{\boldsymbol{\epsilon}}}{\hat{\boldsymbol{\epsilon}}^T \cdot \hat{\boldsymbol{\epsilon}}}, \tag{1}$$

where $\boldsymbol{V}$ denotes the (usually row-standardized) spatial link matrix (cf. Cliff and Ord, 1973).

Anselin proposed the local Moran's $\mathcal{I}_i$ statistic to test for local spatial dependence and to detect spatial objects with influence on global Moran's $\mathcal{I}$ (Anselin, 1995). He states two requirements that have to be fulfilled by a sensible local indicator:

1. The local statistic should give an indication of outstanding clustering in either positive or negative spatially distributed regression residuals, as well as pointing to significant spatial outliers in the regression residuals that do not fit into their surrounding environment. The notion of local clusters is related to positive spatial autocorrelation whereas the notion of spatial outliers is tied to negative spatial autocorrelation.

2. The sum of the local indicators should be proportional to a global indicator of spatial association defined over all spatial objects. Consequently, the global statistic can be broken down into its local components and their impact on the global statistic can be investigated.

For the numerator of the local Moran's $\mathcal{I}_i$, Anselin (1995) proposed a weighted linear combination of the residuals. Tiefelsdorf and Boots (1997) defined the numerator as a quadratic form in regression residuals with a local spatial link matrix $\boldsymbol{V}_i$ which is analytically equivalent to Anselin's definition. The basic idea is to decompose a global spatial binary link matrix $\boldsymbol{G}$ into n local sparse star-shaped structures $\boldsymbol{G}_i$.

$$\boldsymbol{G}_i \equiv \begin{pmatrix} 0 & \cdots & 0 & g_{1i} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & g_{i-1,i} & 0 & \cdots & 0 \\ g_{i1} & \cdots & g_{i,i-1} & 0 & g_{i,i+1} & \cdots & g_{in} \\ 0 & \cdots & 0 & g_{i+1,i} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & g_{ni} & 0 & \cdots & 0 \end{pmatrix}$$

Note that there are two eigenvalues of $\boldsymbol{G}_i$ which are nonzero and the whole spectrum can be given in analytical terms as

$$\left\{ -\sqrt{\sum_{j=1}^{n} g_{ij}^2}, 0, \ldots, 0, \sqrt{\sum_{j=1}^{n} g_{ij}^2} \right\} \tag{2}$$

The observed value of the local Moran's $\mathcal{I}_i$ is then given by

$$\mathcal{I}_i = \frac{\hat{\boldsymbol{\epsilon}}^T \cdot \boldsymbol{V}_i \cdot \hat{\boldsymbol{\epsilon}}}{\hat{\boldsymbol{\epsilon}}^T \cdot \hat{\boldsymbol{\epsilon}}} \tag{3}$$

where $\boldsymbol{V}_i \equiv s_i \cdot \boldsymbol{G}_i$ is the local spatial link matrix with a specific s_i for a given coding scheme (e.g. $s_i = \frac{1}{2 \cdot d_i}$ and $(d_1, ..., d_n)^T = \boldsymbol{G} \cdot \vec{1}$ for row-sum standardization). These definitions assure the additivity requirement, i.e. that the symmetric coded global link matrix is the sum of the coded local link matrices

$$\frac{1}{2}\left(\boldsymbol{V} + \boldsymbol{V}^T\right) = \sum_{i=1}^{n} \boldsymbol{V}_i \tag{4}$$

and the global Moran's $\mathcal{I}$ is the sum of the local Moran's $\mathcal{I}_i$'s.

Local Moran's $\mathcal{I}_i$ have increasingly been used for the spatial detection of so called hotspots or clusters, i.e. significantly deviant local outliers (cf. McLaughlin and Boscoe, 2007). They came to be an almost unrivaled tool for this purpose, for a recent alternative, however, see Ainsworth and Dean (2008).

2 Moran's $\mathcal{I}$-Tests

To decide whether a particular realized value of $\mathcal{I}_i$ is indicative of spatial independence (Null-hypothesis) or dependence (Alternative), the application of a statistical test is required, and thus the distributions of $\mathcal{I}_i$ under both hypotheses, if both size and power is to be evaluated.

The Moran's $\mathcal{I}$ test statistic measures the intensity of spatial autocorrelation in a spatial process but not directly the spatial autocorrelation level ρ (see Li et al., 2007, for an approximate profile-likelihood estimator). Tests for positive correlation are much more relevant in practice, because negative spatial autocorrelation very rarely appears in the real world.

For a properly conducted power analysis we require the ability to calculate the distribution of our test statistic given spatial dependence driven by a particular spatial process. This is similarly so, when we want to detect local outliers against the background of a global process. Although both apply different techniques, much general discussion on the latter issue can be found in Ord and Getis (2001) and Goovaerts and Jacquez (2005) respectively. However, note that power analyses are rarely performed in this context with a short paragraph in Tiefelsdorf (1998) being a rare exception.

2.1 Normal Approximation

One of the most commonly used methods is to employ the asymptotic normal distribution (derived by Cliff and Ord, 1973) as an approximation, and test the standardized value of the statistic

$$z(\mathcal{I}) = \frac{\mathcal{I}_0 - E[\mathcal{I}]}{\sqrt{V[\mathcal{I}]}} \quad \sim \mathcal{N}(0, 1) \tag{5}$$

against a standard normal.

For this purpose and the purpose of power evaluation one needs to evaluate the moments under the assumption of spatial independence as well as under a given spatial process for the alternative.

The expectation and variance under the assumption of spatial independence can be easily derived from the central moments of $\mathcal{I}$ (see Tiefelsdorf, 2000). They yield

$$E\left[\mathcal{I}\,|H_0\right] = \frac{tr\left(\boldsymbol{K}\right)}{n-k} = \overline{\gamma} \tag{6}$$

and

$$V\left[\mathcal{I}\,|H_0\right] = \frac{2\left((n-k)tr(\boldsymbol{K}^2) - tr(\boldsymbol{K})^2\right)}{(n-k)^2\,(n-k+2)} \tag{7}$$

respectively, where $\{\gamma_1, ..., \gamma_{n-k}, 0, ..., 0\}$ are the eigenvalues of the matrix $\boldsymbol{M} \cdot \frac{1}{2} \cdot \left(\boldsymbol{V} + \boldsymbol{V}^T\right) \cdot \boldsymbol{M}$. The sum of the eigenvalues of $\boldsymbol{K} \equiv \boldsymbol{M} \cdot \frac{1}{2} \cdot \left(\boldsymbol{V} + \boldsymbol{V}^T\right) \cdot \boldsymbol{M}$ can be evaluated with the trace operator $tr\left(\boldsymbol{K}^p\right) = \sum_{i=1}^{n-k} \gamma_i^p$.

The moments of Moran's $\mathcal{I}$ under the influence of a spatial process are not so easy to derive. For more details see Tiefelsdorf (2000).

Under the influence of a spatial process the random errors $\boldsymbol{\epsilon}$ are normally distributed with covariance matrix $\sigma^2\boldsymbol{\Omega}$, and hence the regression residuals $\hat{\boldsymbol{\epsilon}}$ are normally distributed with covariance matrix $\sigma^2\boldsymbol{M\Omega M}$. The expectation of the random errors $\boldsymbol{\epsilon}$ and the expectation of the regression residuals $\hat{\boldsymbol{\epsilon}}$ are zero, which is eventually important as it leads to central χ^2-distributed variables. The structure of the matrix $\boldsymbol{\Omega}$ depends upon the spatial process that is assumed to generate the data under the alternative. Throughout the paper we will assume a Gaussian simultaneous spatial autoregressive (SAR) process, i.e. $\boldsymbol{\epsilon} = \rho\boldsymbol{V} \cdot \boldsymbol{\epsilon} + \boldsymbol{\eta}$ with $\boldsymbol{\eta} \sim \mathcal{N}(\boldsymbol{0},\, \sigma^2 \cdot \boldsymbol{I})$, which yields $\boldsymbol{\Omega}^{\frac{1}{2}} = \left(\boldsymbol{I} - \rho\boldsymbol{V}^T\right)^{-1}$, without loss of generality.

Let us then define

$$\begin{aligned}
\boldsymbol{A} \equiv& \boldsymbol{\Omega}^{T\frac{1}{2}} \cdot \boldsymbol{M} \cdot \frac{1}{2} \cdot \left(\boldsymbol{V} + \boldsymbol{V}^T\right) \cdot \boldsymbol{M} \cdot \boldsymbol{\Omega}^{\frac{1}{2}} \\
\boldsymbol{B} \equiv& \boldsymbol{\Omega}^{T\frac{1}{2}} \cdot \boldsymbol{M} \cdot \boldsymbol{\Omega}^{\frac{1}{2}}
\end{aligned}$$

with β_i being the eigenvalues of $\boldsymbol{B}$ and $\boldsymbol{P}$ a $n \times n$ matrix whose columns are the normalized eigenvectors of $\boldsymbol{B}$. Because of the rank defect of the projection matrix $\boldsymbol{M}$ only $n-k$ eigenvalues of $\boldsymbol{B}$ are non zero.

The conditional expectation of Moran's $\mathcal{I}$ is then given by

$$E\left[\mathcal{I} \mid H_1\right] = \int_0^\infty \prod_{i=1}^{n-k} \left(1 + 2 \cdot \beta_i \cdot t\right)^{-\frac{1}{2}} \cdot \sum_{i=1}^{n-k} \frac{h_{ii}}{1 + 2 \cdot \beta_i \cdot t} dt$$

and the conditional second moment of the Moran's $\mathcal{I}$ can be written as

$$E\left[\mathcal{I}^2 \mid H_1\right] = \int_0^\infty \prod_{i=1}^{n-k} \left(1 + 2 \cdot \beta_i \cdot t\right)^{-\frac{1}{2}} \cdot \sum_{i=1}^{n-k} \sum_{j=1}^{n-k} \frac{\left(h_{ii} \cdot h_{jj} + 2 \cdot h_{ij}^2\right) \cdot t}{\left(1 + 2 \cdot \beta_i \cdot t\right) \cdot \left(1 + 2 \cdot \beta_j \cdot t\right)} dt$$

where the h_{ij} are the diagonal elements of the matrix $\boldsymbol{P}^T \cdot \boldsymbol{A} \cdot \boldsymbol{P}$.

An alternative to the normal approximation, which is implemented in some software packages (GeoDa) is a permutation based approach usually applied directly to the variable under analysis. There are asymptotic arguments for extending this practice also to regression residuals - see Jacqmin-Gadda et al. (1997) in this context following the reasoning of Schmoyer (1994). However, note that the permutations are merely employed to generate the first two moments, and a normal approximation is used for evaluating the whole distribution anyway. Thus for power analysis purposes the two approaches differ little.

2.2 The exact distribution

It turns out that the normal approximation is not very suitable for small lattices (i.e. especially the case of local Moran's $\mathcal{I}_i$), so that other approximation methods were suggested, e.g. the Saddlepoint approximation by Tiefelsdorf (2002). However, due to increasing computing power, it has become possible to evaluate the numerically demanding exact distribution of the Moran's $\mathcal{I}$ statistic in shorter time for even big lattices. We will furthermore show that especially for power computations the errors induced by the using the normal approximation can be severe.

In the following, we will give a brief review of the theory essentially following the derivation from Tiefelsdorf and Boots (1995) (an independent development can be found in Hepple (1998)). A comprehensive exposition of this derivation and many related issues including some comments on power analyses can be found in the monograph by Tiefelsdorf (2000).

Under the influence of a spatial process, the conditional distribution of Moran's $\mathcal{I}$ given the observed value $\mathcal{I}_0$ and a hypothetical spatial process generating $\sigma^2\boldsymbol{\Omega}$ can be written as

$$F\left(\mathcal{I}_0\,|H_1\right) = P\left(\frac{\boldsymbol{\delta}^T\cdot\boldsymbol{\Omega}^{T\frac{1}{2}}\cdot\boldsymbol{M}\cdot\frac{1}{2}\cdot\left(\boldsymbol{V}+\boldsymbol{V}^T\right)\cdot\boldsymbol{M}\cdot\boldsymbol{\Omega}^{\frac{1}{2}}\cdot\boldsymbol{\delta}}{\boldsymbol{\delta}^T\cdot\boldsymbol{\Omega}^{T\frac{1}{2}}\cdot\boldsymbol{M}\cdot\boldsymbol{\Omega}^{\frac{1}{2}}\cdot\boldsymbol{\delta}}\leq\mathcal{I}_0\right) \tag{8}$$

$$= P\left(\boldsymbol{\delta}^T\cdot\boldsymbol{\Omega}^{T\frac{1}{2}}\cdot\boldsymbol{M}\cdot\left[\frac{1}{2}\cdot\left(\boldsymbol{V}+\boldsymbol{V}^T\right)-\mathcal{I}_0\cdot\boldsymbol{I}\right]\cdot\boldsymbol{M}\cdot\boldsymbol{\Omega}^{\frac{1}{2}}\cdot\boldsymbol{\delta}\leq 0\right) \tag{9}$$

By the spectral decomposition Theorem

$$\boldsymbol{L}_{H1}\equiv\boldsymbol{\Omega}^{T\frac{1}{2}}\cdot\boldsymbol{M}\cdot\left[\frac{1}{2}\cdot\left(\boldsymbol{V}+\boldsymbol{V}^T\right)-\mathcal{I}_0\cdot\boldsymbol{I}\right]\cdot\boldsymbol{M}\cdot\boldsymbol{\Omega}^{\frac{1}{2}} \tag{10}$$

(note that $\boldsymbol{L}_{H1}$ is symmetric) can be written as $\boldsymbol{L}_{H1}=\boldsymbol{H}^T\cdot\boldsymbol{\Lambda}\cdot\boldsymbol{H}$, where $\boldsymbol{H}$ is the matrix of the normalized eigenvectors and $\boldsymbol{\Lambda}=diag(\lambda_1,...\lambda_n)$ is the diagonal eigenvalue matrix of $\boldsymbol{L}_{H1}$ given in equation (10). Substituting into equation (9) we get

$$F\left(\mathcal{I}_0\,|H_1\right) = P\left(\boldsymbol{\delta}^T\cdot\boldsymbol{H}^T\cdot\boldsymbol{\Lambda}\cdot\boldsymbol{H}\cdot\boldsymbol{\delta}\leq 0\,|H_1\right)$$

Because the random error vector $\boldsymbol{\delta}$ belongs to the class of the spherically symmetric distributions, the orthogonal transformation $\boldsymbol{\eta}\equiv\boldsymbol{H}\cdot\boldsymbol{\delta}$ is again independent normal distributed with $\boldsymbol{\eta}\sim\mathcal{N}\left(\boldsymbol{0},\boldsymbol{I}\right)$ (Tiefelsdorf, 2000).

So the conditional distribution of Moran's $\mathcal{I}$ is given by

$$F\left(\mathcal{I}_0\,|H_1\right) = P\left(\sum_{i=1}^{n} \lambda_i \cdot \eta_i^2 \leq 0\,|H_1\right), \tag{11}$$

and this enables us to use Imhof's formula because $\sum_{i=1}^{n} \lambda_i \cdot \eta_i^2$ is a weighted sum of χ_1^2-distributed variables.

> **(Imhof's Formula)** (Imhof, 1961) The distribution function $F(y)$ of the weighted sum of independent central χ^2-distributed variables is given by
>
> $$F(y) = P(Y \leq y) = \frac{1}{2} - \frac{1}{\pi} \int_0^{\infty} \frac{1}{u} \cdot sin\left(\boldsymbol{\Theta}(u)\right) \cdot \xi(u)\, du$$
>
> Where $X_1, X_2, \ldots, X_n$ are independent χ_1^2-distributed random variables, with the weights $\lambda_1, \lambda_2 \ldots, \lambda_n \in \mathbb{R}$. Thus the weighted sum $Y = \lambda_1 \cdot X_1 + \lambda_2 \cdot X_2 + \ldots + \lambda_n \cdot X_n$.
>
> The two functions $\boldsymbol{\Theta}(u)$ and $\xi(u)$ are given by
>
> $$\boldsymbol{\Theta}(u) = \frac{1}{2} \sum_{j=1}^{n} arctan(u\lambda_j) - \frac{1}{2} u\, y$$
>
> $$\xi(u) = \prod_{j=1}^{n} (1 + u^2 \lambda_j^2)^{-\frac{1}{4}}$$

Note that all zero eigenvalues can be ignored, and because of $y = 0$, the term $-\frac{1}{2} u \cdot y$ in Imhof's formula is irrelevant for our purposes.

Another way is the direct evaluation of the complex-valued characteristic function of a weighted sum of χ^2-distributed variables. It has not succeeded in practice, because the calculation is not easy to implement and the approach above with real-valued integration is much easier to handle. Here, the solution of the integral in Imhof's formula can be approximated by numerical integration. The behavior of the improper integral at $u = 0$ and at $u \to \infty$ have to be considered especially, yielding starting and truncation values respectively (see Tiefelsdorf, 2000, for details).

Under the assumption of spatial independence ($\rho = 0$) the covariance matrix $\sigma^2 \boldsymbol{\Omega}$ reduces to $\sigma^2 \boldsymbol{I}$ and we are able to simplify the conditional distribution of Moran's $\mathcal{I}$ given by equation (9) to

$$F\left(\mathcal{I}_0\,|H_0\right) = P\left(\boldsymbol{\delta}^T \cdot \boldsymbol{M} \cdot \left[\frac{1}{2} \cdot \left(\boldsymbol{V} + \boldsymbol{V}^T\right) - \mathcal{I}_0 \cdot \boldsymbol{I}\right] \cdot \boldsymbol{M} \cdot \boldsymbol{\delta} \leq 0\right) \tag{12}$$

with

$$\boldsymbol{L}_{H0} \equiv \boldsymbol{M} \cdot \left[\frac{1}{2} \cdot \left(\boldsymbol{V} + \boldsymbol{V}^T\right) - \mathcal{I}_0 \cdot \boldsymbol{I}\right] \cdot \boldsymbol{M} \tag{13}$$

and similar calculations as before give us

$$F\left(\mathcal{I}_0\,|H_0\right) = P\left(\sum_{i=1}^{n} \lambda_i \cdot \eta_i^2 \leq 0\,|H_0\right) \quad = P\left(\sum_{i=1}^{n-k} \left(\gamma_i - \mathcal{I}_0\right) \cdot \eta_i^2 \leq 0\,|H_0\right) \qquad (14)$$

where $\boldsymbol{L}_{H0} = \boldsymbol{H}^T \cdot \boldsymbol{\Lambda} \cdot \boldsymbol{H}$, $\boldsymbol{H}$ is the matrix of the normalized eigenvectors and $\boldsymbol{\Lambda} = diag(\lambda_1, ...\lambda_n)$ is the diagonal eigenvalue matrix of $\boldsymbol{L}_{H0}$ given in equation (13).

The latter equality, where $\boldsymbol{\Gamma} = diag(\gamma_1, \ldots, \gamma_{n-k}, 0, \ldots, 0)$ is the diagonal matrix of the eigenvalues of the matrix $\boldsymbol{M} \cdot \frac{1}{2} \cdot \left(\boldsymbol{V} + \boldsymbol{V}^T\right) \cdot \boldsymbol{M}$ is a simplification allowing the quick calculation for each observed $\mathcal{I}_0$.

Again we are able to use Imhof's formula because $\sum_{i=1}^{n-k} \left(\gamma_i - \mathcal{I}_0\right) \eta_i^2$ is a weighted sum of χ_1^2-distributed variables. Comparing with equation (14) we get that $\lambda_i = \gamma_i - \mathcal{I}_0$ for $i = \{1, ..., n-k\}$ and $\lambda_j = 0$ for the remaining $j = \{n-k+1, ..., n\}$.

It is useful to know that the feasible range of the exact distribution of Moran's $\mathcal{I}$ is given by $[\gamma_1, \gamma_{n-k}]$, i.e. the interval over which we have nonzero probability (cf. Dhrymes, 1984).

All of the above also holds for computing size and power for local $\mathcal{I}_i$ tests. However, here we face the additional issue of size inflation by multiple testing. And although we will employ the common Bonferroni correction based on n throughout the paper, it is evident that through the correlated nature of the data/residuals this leads to being overconservative. It may be reasonable to base the correction on only a fraction of n representing the equivalent number of independent observations. A detailed discussion of many testing issues (except power analysis) for local spatial indicators can be found in Leung et al. (2003).

3 Implementation Issues

The initial implementation was undertaken by Markus Reder, using the R open source statistical computation environment (R Development Core Team, 2007). An integration of this work into the **spdep** package by (Bivand et al., 2008) has been released. In this section, we will discuss how the implementation has been carried out.

The **spdep** package had contained a number of global and local versions of the Moran's
mi, including versions using the Saddlepoint approximation contributed by Michael Tiefelsdorf, and implementing his presentation of the approach in Tiefelsdorf (2002). In the global case, eigenvalues have to be found for an $n \times n$ matrix. This also applies in the local case when testing under the alternative that spatial dependence is present, which we will need to do for the power analysis.

Things can, however, be simplified for evaluating the exact distributions for local Moran's
mi under the null hypothesis, in the same way as for the Saddlepoint approximation. Here, under the assumption of spatial independence, we have to derive the eigenvalues of the matrix $\boldsymbol{M} \cdot \boldsymbol{V}_i \cdot \boldsymbol{M}$. However, as mentioned above, only two eigenvalues of

the local spatial link matrix $\boldsymbol{V}_i$ are nonzero, thus only two eigenvalues of $\boldsymbol{M} \cdot \boldsymbol{V}_i \cdot \boldsymbol{M}$ are nonzero. It is possible to derive these two eigenvalues directly by the following equations (Tiefelsdorf, 2002)

$$\gamma_{1,n-k} = \frac{1}{2} \cdot \left(t_1 \pm \sqrt{2 \cdot t_2 - t_1^2} \right) \tag{15}$$

where

$$\begin{aligned} t_1 &= tr(\boldsymbol{M} \cdot \boldsymbol{V}_i \cdot \boldsymbol{M}) \\ t_2 &= tr\left((\boldsymbol{M} \cdot \boldsymbol{V}_i \cdot \boldsymbol{M})^2 \right) \end{aligned}$$

A further simplification is to use the following identities.

$$\begin{aligned} tr(\boldsymbol{M} \cdot \boldsymbol{V}_i \cdot \boldsymbol{M}) =& - tr\left(\boldsymbol{X}^T \boldsymbol{V}_i \boldsymbol{X} \left(\boldsymbol{X}^T \boldsymbol{X} \right)^{-1} \right) \\ tr\left((\boldsymbol{M} \cdot \boldsymbol{V}_i \cdot \boldsymbol{M})^2 \right) =& tr\left(\boldsymbol{V}_i^2 \right) - tr\left(\boldsymbol{X}^T \boldsymbol{V}_i^2 \boldsymbol{X} \left(\boldsymbol{X}^T \boldsymbol{X} \right)^{-1} \right) \\ & + tr\left(\left(\boldsymbol{X}^T \boldsymbol{V}_i \boldsymbol{X} \left(\boldsymbol{X}^T \boldsymbol{X} \right)^{-1} \right)^2 \right) \end{aligned}$$

Because most matrix products on the right-hand side involve only $k \times k$ matrices instead of $n \times n$ matrices, the required number of computation operations are reduced substantially.

Since this approach had been implemented in the `localmoran.sad` function for the Saddlepoint approximation, introducing a corresponding `localmoran.exact` function was not difficult. Both functions take care to use the residuals of a linear model as their first argument, thus implementing univariate local Moran's I as an analysis of the residuals of a model including only the intercept. If explanatory variables can be added, thus removing mis-specification, testing the residuals will be more robust. In the substantive example below, we will only use an intercept-only model, although it may well be that using explanatory variables and/or weights to reduce possible heteroskedasticity would have permitted us to draw clearer conclusions.

Power analysis involves working under the alternative hypothesis in order to establish whether we reject the null only when we should, and not otherwise. The key issue raised in work on exact inference for Moran's I is that, for the global case, the normal approximation may lead to over-eager rejection of the null hypothesis. In the local case, this is much more likely to happen, not least because of the very limited number of observations involved in each inference.

The implementation of the functions has been conducted so that the intermediate objects, matrices, eigenvalues, and local sparse star-shaped structures can optionally be returned. This avoids their sometimes costly re-computation during analysis. The internal code for computing the exact probability values has been split out within the **spdep** namespace, so that it can be used with the intermediate objects in order to conduct power analyses. These internal functions are not exported, but can be accessed where necessary. While the substantive cases will be presented in the next section, we will show how intermediate objects and internal functions can be used, here to produce some of the values shown in Table 1.

First we will read in a shapefile containing the boundaries of communities in Upper Austria (see Figure 1, panel c). A list of neighbours based on shared boundary points is constructed using the `poly2nb` function, and finally the net migration rate per 1000 inhabitants for 2002 is added to the spatial polygon data frame object:

```
> library(maptools)
> UA <- readShapePoly("gemeinden.shp")
> library(spdep)
> UA_nb <- poly2nb(UA)
> load("mig2002.rDATA")
> UA$mig2002 <- mig2002
```

We have no other model of the data than the intercept, so fit a null linear model. We use the community of Atzesberg as our single case here, for reasons explained in the following section. The `localmoran.exact` takes the object containing the null model fit as its first argument, the selected community as a single element vector (if the argument is omitted, all communities will be chosen), and the list of coniguous neighbours and the `"W"` style argument for row-standardisation:

```
> lmobj <- lm(mig2002 ~ 1, UA)
> Atzesberg <- which(UA$GEM_NAME == "Atzesberg")
> reso_Atzesberg <- localmoran.exact(lmobj, select = Atzesberg,
+     nb = UA_nb, style = "W")
> reso_Atzesberg

      Local Morans I Exact SD  Pr. (exact)
36 35       5.717311 5.085056 1.837589e-07
```

The object returned is a single member list of `"htest"` objects, single member because we only selected one community — what is displayed is the output of the `print` method for this class. Each object in the list also has some intermediate information returned with it, here the pair of analytical eigenvalues for the product of the star-shaped weights matrix for Atzesberg and the projection matrix from `lmobj`, computed as shown above.

```
> Gamma <- reso_Atzesberg[[1]]$gamma
> Gamma

[1] -99.3332  98.3332
```

We also prepare some extra arguments for the next step, including the Bonferroni-adjusted 0.05 probability value; our scenario is that we have actually computed local Moran's $\mathcal{I}_i$ for all the communities in Upper Austria, hence the adjustment:

```
> N <- length(UA_nb)
> np2 <- N - (2 + lmobj$rank)
> pp <- 0.05/N
> pp

[1] 0.0001123596
```

Our aim is to find the values of Moran's $\mathcal{I}_i$ for the Atzesberg star-shaped weights matrix and the projection matrix from `lmobj` corresponding to a Bonferroni-adjusted 0.05 probability value. We write a function to hand to `optimize` to do a line search within the limits given, reporting the value at the function minimum. The function

calls an internal function within the **spdep** namespace to return the exact probability value for varying values `x` of Moran's $\mathcal{I}_i$, given the eigenvalues, that this is a local Moran's $\mathcal{I}_i$, and the value of `np2` for this data set:

```
> f <- function(x) abs(pp - spdep:::exactMoran(x, Gamma, type = "Local",
+     np2 = np2)$p.value)
> I_ex <- optimize(f, interval = c(-5, 5))$minimum
> I_ex
```

```
[1] 3.112458
```

If we calculate the first two moments using the eigenvalues already made available, we can use `qnorm` to find the value of local Moran's $\mathcal{I}_i$ corresponding to a Bonferroni-adjusted 0.05 probability value for the Normal approximation and for our specific setting:

```
> gamma <- c(Gamma[1], rep(0, np2), Gamma[2])
> mu1 <- mean(gamma)
> mu2 <- 2/(N - 1 + 2) * mean((gamma - mu1)^2)
> I_nv <- qnorm(pp, mean = mu1, sd = sqrt(mu2), lower.tail = FALSE)
> I_nv
```

```
[1] 1.636615
```

We see that it is very much smaller than the value of local Moran's $\mathcal{I}_i$ found using the exact Bonferroni-adjusted 0.05 probability value. Cross-checking, we can compute the exact probability value of the value of local Moran's $\mathcal{I}_i$ corresponding to a Bonferroni-adjusted 0.05 probability value:

```
> p_ex <- spdep:::exactMoran(I_nv, Gamma, type = "Local", np2 = np2)$p.value
> p_ex
```

```
[1] 0.004419898
```

Going the other way, we can find the Normal approximation probability value for the value of local Moran's $\mathcal{I}_i$ corresponding to a Bonferroni-adjusted 0.05 probability value:

```
> p_nv <- pnorm(I_ex, mean = mu1, sd = sqrt(mu2), lower.tail = FALSE)
> p_nv
```

```
[1] 1.175178e-12
```

These code snippets simply illustrate how the conbination of returning intermediate objects, together with modularising internal code in functional form within the user-visible exact Moran functions provided in the **spdep** package, provides the entry points needed to make power analysis feasible. The complete script for reproducing the substantive cases is available from the authors and from the journal electronic supplementary material site, and may be used with released versions of **spdep**.

In performance terms, the comments made by Tiefelsdorf (2002) continue to hold, but, given advances in memory capacity and processor speed, the advantage of the Saddlepoint approximation over the exact distribution for the null case, assuming the absence of spatial autocorrelation, is less pronounced. In our case selecting all 445 Upper Austrian communities (by omitting the `select=` argument), compared to 219 counties in the former GDR in Tiefelsdorf (2002, p. 203), we see that (for a 2001-built 1.5GHz i386 platform with 1GB memory):

```
> system.time(reso_Atzesberg_E <- localmoran.exact(lmobj, nb = UA_nb,
+     style = "W"))

   user  system elapsed
 14.564   0.009  14.695

> system.time(reso_Atzesberg_S <- localmoran.sad(lmobj, nb = UA_nb,
+     style = "W"))

   user  system elapsed
 13.340   0.005  13.348

> Omega <- diag(length(UA_nb))
> system.time(reso_Atzesberg_Ea <- localmoran.exact.alt(lmobj,
+     nb = UA_nb, style = "W", Omega = Omega))

   user  system elapsed
593.112  37.404 631.015
```

Naturally, when the analytical method discussed above cannot be used, the numerical function call count for large matrices is increased radically, so that the results for the null exact distribution in the first case running `localmoran.exact` and the third running `localmoran.exact.alt` with $\mathbf{\Omega}$ set to the identity matrix are identical, but with a very large difference in timings. A final problem in power analysis is that the `spdep:::H1_moments` internal function for generating moments for the normal approximation under the alternative hypothesis continues to be very compute-intensive. Note, however, that local Moran's $\mathcal{I}_i$ under the alternative for many observations is embarassingly parallelizable, although here only a single processor core has been used. Two processes running with two different `select=` argument vector values would effectively halve the processing time.

4 Substantive Cases

We intend to demonstrate the potential drawbacks of using the normal approximation through power analyses on two artificial and one real example. We will use three different areas with the row-standardized coding scheme: a regular *5×5 grid* (Figure 1 a), the *B07-structure* (Figure 1 b) from the set of fourteen maximally connected planar spatial structures called the B-series with a fixed number $n = 8$ nodes and an overall connectivity $D = 36$ from Boots and Royle (1991) and *Upper Austria* (Figure 1 c), a federal state of Austria with 445 communities.

4.1 Artificial Examples

4.1.1 The 5×5 grid

In Figure 2, left panel, the deviation between the normal approximation and the exact distribution under the Null hypothesis is plotted against global Moran's $\mathcal{I}$. One may note that there is a reasonable agreement over a wide range of values. In particular, for the value 0.20545, for which we are just able to reject the Null for a one-sided size 0.05 test, which is represented by the dashed line (and any value large than that), the exact p-value would typically be just less than 0.05. Thus in this

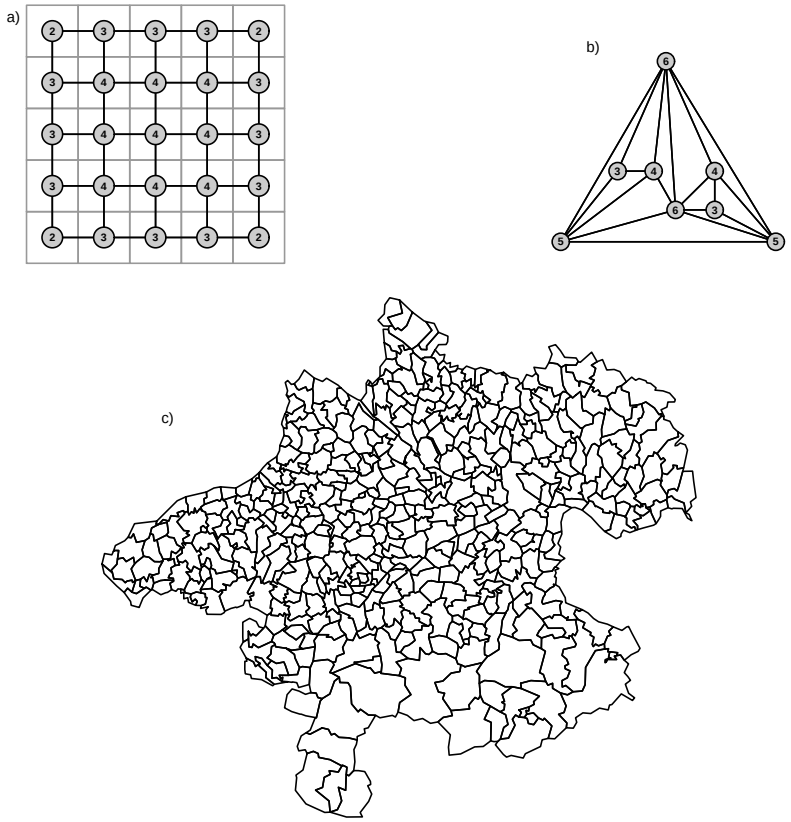

Figure 1: Spatial structures: a) 5×5 grid; b) B07-structure; c) Upper Austria — communities.

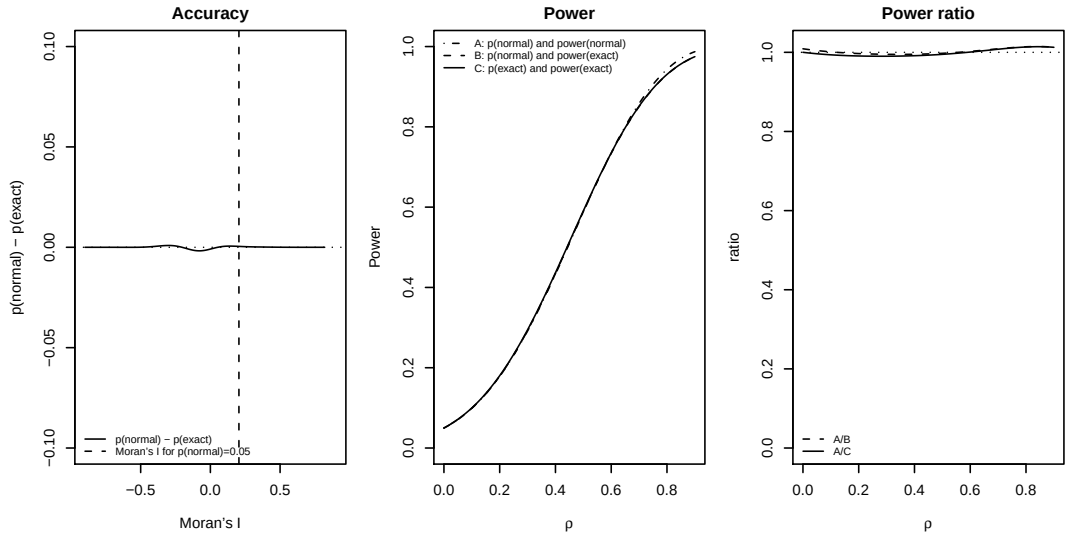

Figure 2: The 5×5 grid — left panel: accuracy for different values of Moran's $\mathcal{I}$; centre panel: power for different values of ρ; right panel: power ratios for different values of ρ.

example the normal approximation does not do any harm, other than being slightly too conservative.

Let us turn to a corresponding power analysis: Figure 2, centre panel, shows three different power functions.

- **A:** Deriving the p-value with the normal approximation hence the critical Moran's $\mathcal{I}$ is 0.20545. Deriving the power for different autocorrelation levels ρ with the normal approximation.

- **B:** Deriving the p-value with the normal approximation hence the critical Moran's $\mathcal{I}$ is 0.20545. Deriving the power for different autocorrelation levels ρ with the exact distribution.

- **C:** Deriving the exact p-value hence the critical Moran's $\mathcal{I}$ is 0.20481. Deriving the power for different autocorrelation levels ρ with the exact distribution.

Note that here the differences are marginal. This can be clearly seen by displaying the ratios of the power functions displayed in Figure 2, right panel, which are close to 1 throughout the range.

4.1.2 B07-structure

The picture is very different for our second example. We are adding equivalent results for the Saddlepoint approximation accuracy and power analyses, to match those in Tiefelsdorf (2002) for this lattice.

Figure 3, upper panel, again plots the deviation between the normal approximation and the exact approach versus Moran's $\mathcal{I}$. Now for the normal approximation critical value of 0.092283 (and all larger values) the exact p-value is larger than 0.05, thus we would too quickly reject the Null hypothesis; the exact critical value is 0.13850 here. The differences between the normal and Saddlepoint approximations and the normal approximation and exact values are very similar, bearing out the conclusions in Tiefelsdorf (2002).

Note that this mistake may manifest itself also in the power analysis. Figure 3 centre panel now shows six different power functions.

- **A:** Normal approximation test, hence the critical Moran's $\mathcal{I}$ is 0.092283. Deriving the power for different autocorrelation levels ρ with the normal approximation.

- **B:** Normal approximation test, hence the critical Moran's $\mathcal{I}$ is 0.092283. Deriving the power for different autocorrelation levels ρ with the exact distribution.

- **C:** Exact distribution based test, hence the critical Moran's $\mathcal{I}$ is 0.13850. Deriving the power for different autocorrelation levels ρ with the exact distribution.

- **D:** Normal approximation test, hence the critical Moran's $\mathcal{I}$ is 0.092283. Deriving the power for different autocorrelation levels ρ with the Saddlepoint approximation.

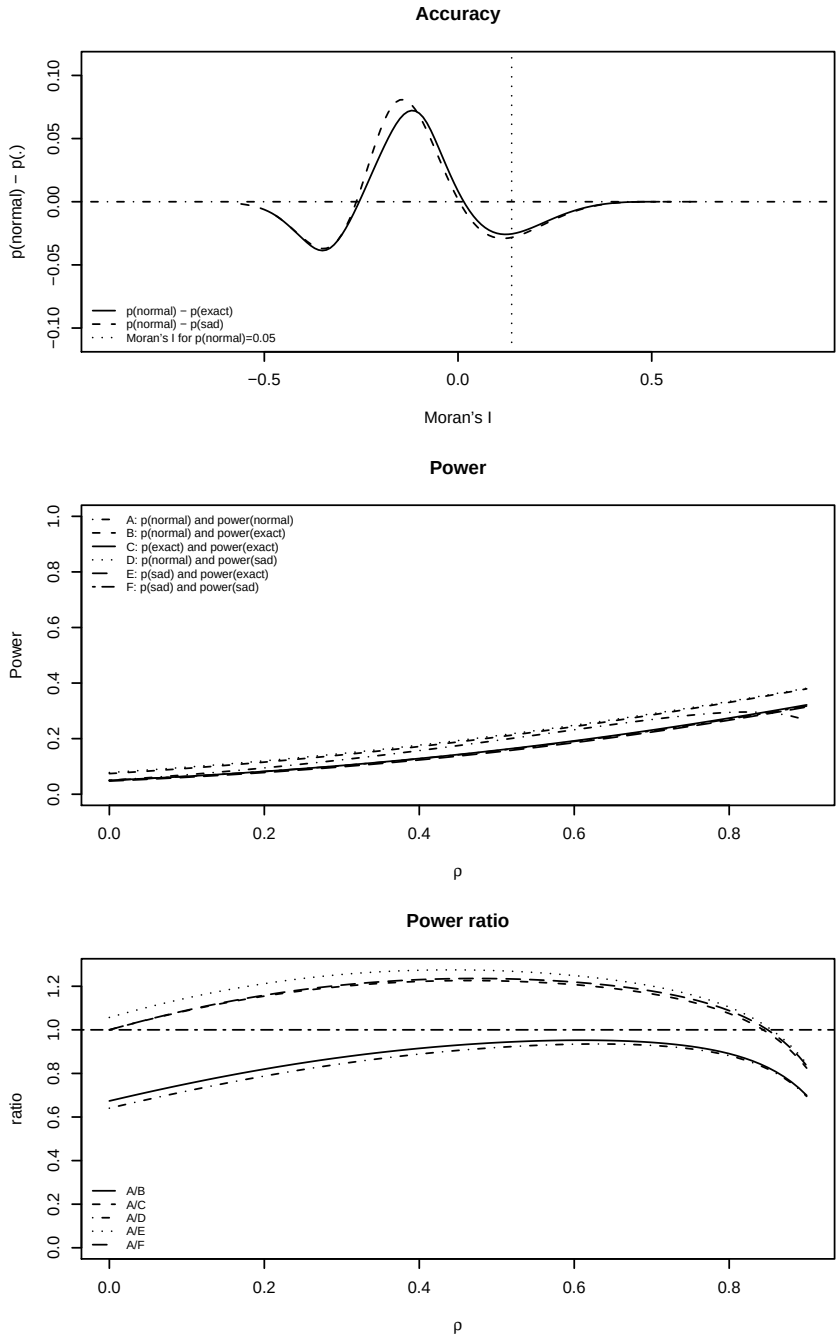

Figure 3: The B07 lattice — upper panel: accuracy for different values of Moran's $\mathcal{I}$; centre panel: power for different values of ρ; lower panel: power ratios for different values of ρ.

- **E:** Saddlepoint approximation based test, hence the critical Moran's $\mathcal{I}$ is 0.14464. Deriving the power for different autocorrelation levels ρ with the exact distribution.

- **F:** Saddlepoint approximation based test, hence the critical Moran's $\mathcal{I}$ is 0.14464. Deriving the power for different autocorrelation levels ρ with the Saddlepoint approximation.

The differences in the reported powers can now be substantial. Note that the power using the normal approximation may go down to 60% (solid A/B line in Figure 3, lower panel) while it could be wrongly reported 20% too high (dashed A/C line). The lines form two groups, with the A/B and A/D lines lying close together — the Normal approximation test with power derived from the exact distribution and the Saddlepoint approximation respectively, and A/C, A/E, and A/F — the Saddlepoint approximation and Exact distribution based tests with power derived from the exact distribution (A/C, A/E) and the Saddlepoint approximation respectively (A/F). Although this example represents a global analysis, it has merits in our context as well because of the small numbers of observations. In addition, we note that choosing the exact distribution based test or the Saddlepoint approximation based test for this lattice configuration is clearly superior to the normal approximation.

4.2 Real Example

Let us now turn to a real local analysis. Figure 4 shows a LISA (local indicators of spatial association) plot for the residuals of a null SAR model, with a global autoregressive coefficient value of 0.1563371, of net migration per 1000 inhabitants for Upper Austria in the year 2002; this global spatial process is significant with a likelihood ratio test value of 0.0312. Additionally, adjusted for the presence of this global process, we find significant local spatial autocorrelation for the *Atzesberg* community, as we can see from the LISA plot.

Figure 5 shows the conditional distribution plot for net migration per 1000 inhabitants for Upper Austria in 2002. Therefore we will derive the conditional distribution function of the local Moran's $\mathcal{I}_i$ $F(\mathcal{I}_i \mid \mathbf{\Omega}(\rho_0))$ adjusted to the force of a global spatial SAR reference process $\mathbf{\Omega}(\rho_0)$. A full conditional distribution map following the suggestions in Tiefelsdorf (2000) is displayed in 5, using the residuals of the null SAR model described above, and the projection matrix from the least squares stage of the GLS SAR fit once ρ_0 has been found by maximum likelihood.

The conditional distribution map should be viewed in conjunction with the empirical cumulative distribution plot of the underlying variable, shown in Figure 6. We see clearly that very many of the communities identified as clusters or hotspots in this example have unusually large or small net migration rates. As is well known, extreme rates can be associated with small populations; this is indeed the case here. The community of Atzesberg had in 2001 only 526 inhabitants, and a net migration rate of -52 per thousand, and two of its five neighbours, Oberkappel and Hörbich, also had very small numbers of inhabitants, and extreme net migration rates (-37 and -25 per thousand). The fourth member of the "cluster", Sarleinsbach, is somewhat larger, with a 2001 population of 2364, and a net migration rate in 2002 of -15 per thousand. Arguably, in a real application, we ought to handle the probable

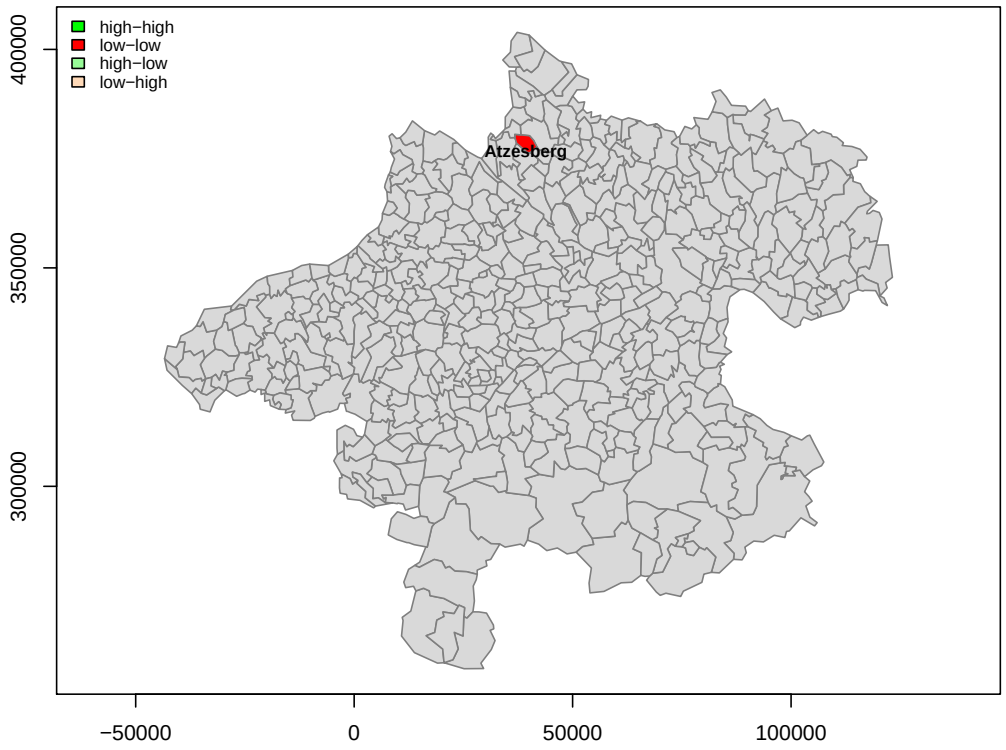

Figure 4: LISA plot of exact local Moran's $\mathcal{I}_i$ p-values, significant at 0.05 after Bonferroni adjustment; null SAR model residuals for net migration per 1000 inhabitants in 2002, Upper Austria.

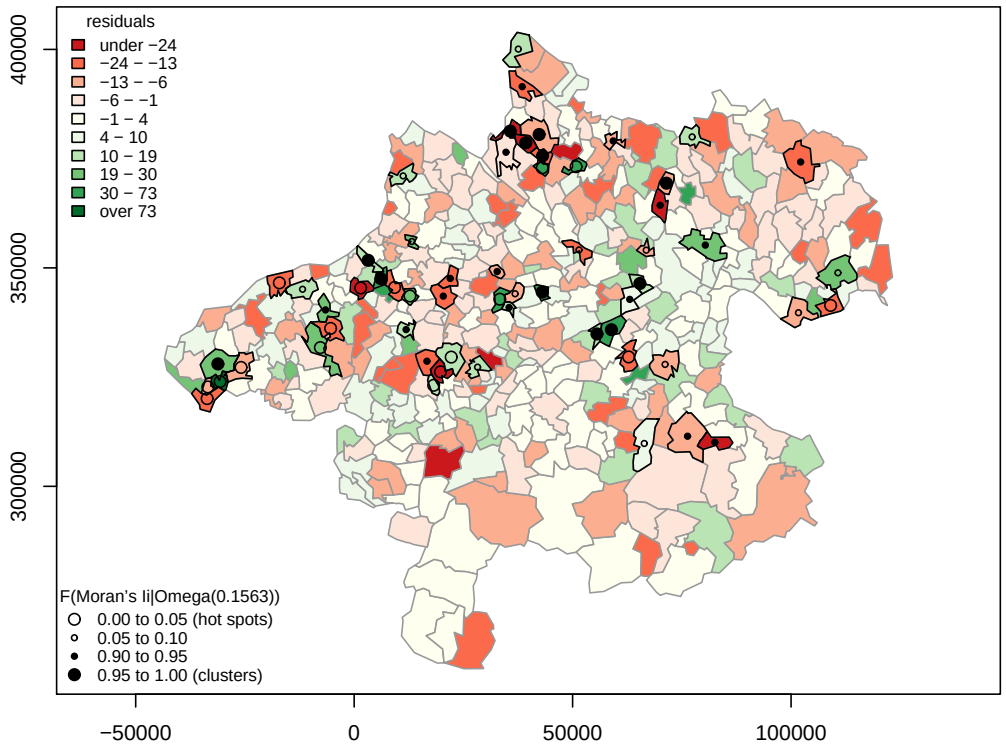

Figure 5: Conditional distribution map of null SAR model residuals for net migration per 1000 inhabitants in 2002, Upper Austria, marked with clusters and hotspots using unadjusted local Moran's $\mathcal{I}_i$ test p-values.

presence of variability in the variance of the rates in our model, before analysing the residuals for local spatial autocorrelation. So although this can be represented as a problem of outliers, it may be more prudent to correct the underlying model first before drawing conclusions. For the purposes of this power analysis, however, we will proceed as though the value reported for Atzesberg can be relied on fully.

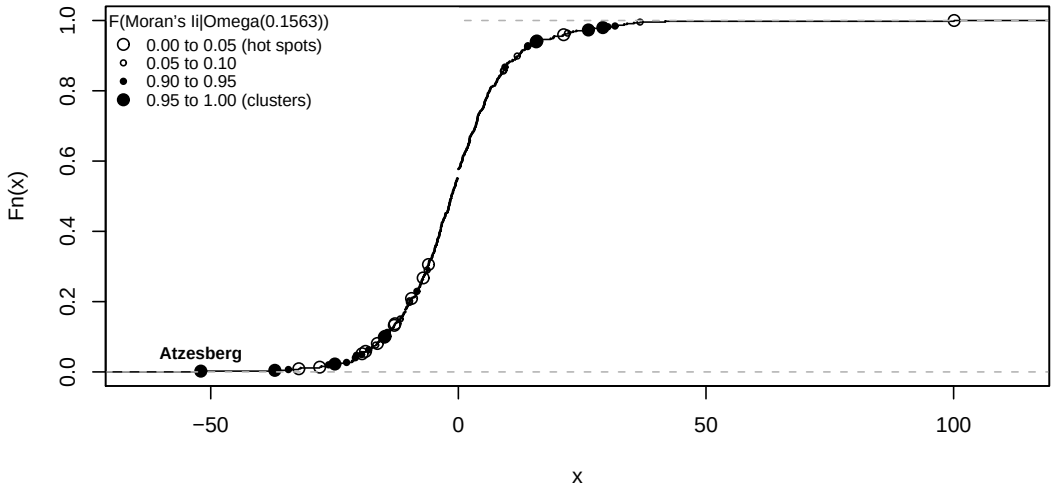

Figure 6: Empirical cumulative distribution function of the null SAR model residuals, marked with clusters and hotspots using unadjusted local Moran's $\mathcal{I}_i$ test p-values.

Table 1: Contrasted values of local Moran's $\mathcal{I}_i$ for Atzesberg for the null SAR model residuals for the Bonferroni adjusted 0.05 level.

	p(normal)	p(sad)	p(exact)	$\mathcal{I}_i$
$\mathcal{I}_i$ normal	0.000112	0.004697	0.004420	1.63661
$\mathcal{I}_i$ sad	7.15e-13	0.000112	0.000104	3.14321
$\mathcal{I}_i$ exact	1.18e-12	0.000121	0.000112	3.11246

We will now investigate more closely the situation of Atzesberg. The community has five neighbours, so the power analysis will apply to other communities with five neighbours in a lattice of this size, because the star-shaped matrices for five neighbours for binary contiguous neighbours and row-standardization will be similar, and the projection matrix only contains the impact of the intercept.

Table 1 shows how the three different tests represent the Bonferroni adjusted 0.05 level, for Atzesberg for the null SAR model residuals, that is 0.000112. The rows of the table show the local Moran's $\mathcal{I}_i$ being tested, and the first three columns the way that the p-value is calculated. The fourth column shows the value of local Moran's $\mathcal{I}_i$ that corresponds to the Bonferroni adjusted 0.05 level for its own test. So the p-values on the diagonal of the left square table are the reference point — if the off-diagonal values differ, they indicate divergence. There is effectively no difference between the exact distribution and Saddlepoint approximation p-values,

but the normal approximation local Moran's $\mathcal{I}_i$ value clearly corresponds to a much larger p-value than the Bonferroni adjusted 0.05 level. The values given here differ slightly than those in section 3, where the residuals were taken from the null linear model rather than the null SAR model.

The upper panel of Figure 7 again displays the deviation between the normal approximation and the exact distribution for its local Moran's $\mathcal{I}_i$, and the deviation between the normal and Saddlepoint approximations. We are using the Bonferroni correction here, and we see that using the normal approximation we would falsely reject the null hypothesis. The sequence of values of local Moran's $\mathcal{I}_i$ used to analyse accuracy are large, because the individual weights in the star-shaped local weights matrices take large values. In this case, with five neighbours, there are only ten weights, summing to 445.

For the same reason, the range of ρ values used in the power analysis seems very small, but simply reflects the feasible range for this star-shaped matrix. Performing a power analysis as above for the B07 lattice, a similar if not more dramatic picture arises.

We see from the centre panel in Figure 7 that the power remains quite low for small to moderate local autocorrelation, but increases for more marked local autocorrelation. We see that the normal approximation (line A) is "over-eager" and rises faster with increasing ρ in the alternative hypothesis than the C, E and F curves, which have very similar outcomes — they take the p-values of the Saddlepoint approximation and exact distribution based tests for local Moran's $\mathcal{I}_i$ at the Bonferroni adjusted 0.05 level.

The middle group of curves are for B and D, where the p-values for the normal approximation local Moran's $\mathcal{I}_i$ at the Bonferroni adjusted 0.05 level are run against the Saddlepoint approximation and exact distribution based powers for increasing values of ρ in the alternative hypothesis.

As we can see from the lower panel of Figure 7, the ratios of these power curves:

- **A:** Normal approximation test, hence the critical Moran's $\mathcal{I}_i$ is 1.6366. Power function with normal approximation.
- **B:** Normal approximation test, hence the critical Moran's $\mathcal{I}_i$ is 1.6366. Exact distribution power function.
- **C:** Exact distribution test, hence the critical Moran's $\mathcal{I}_i$ is 3.1125. Exact distribution power function.
- **D:** Normal approximation test, hence the critical Moran's $\mathcal{I}_i$ is 1.6366. Saddlepoint approximation power function.
- **E:** Saddlepoint approximation based test, hence the critical Moran's $\mathcal{I}_i$ is 3.1432. Exact distribution power function.
- **F:** Saddlepoint approximation based test, hence the critical Moran's $\mathcal{I}_i$ is 3.1432. Saddlepoint approximation power function.

exhibit more extreme behavior than in the B07 lattice case. Here the power could be reported four times higher using the normal approximation than its exact value, for

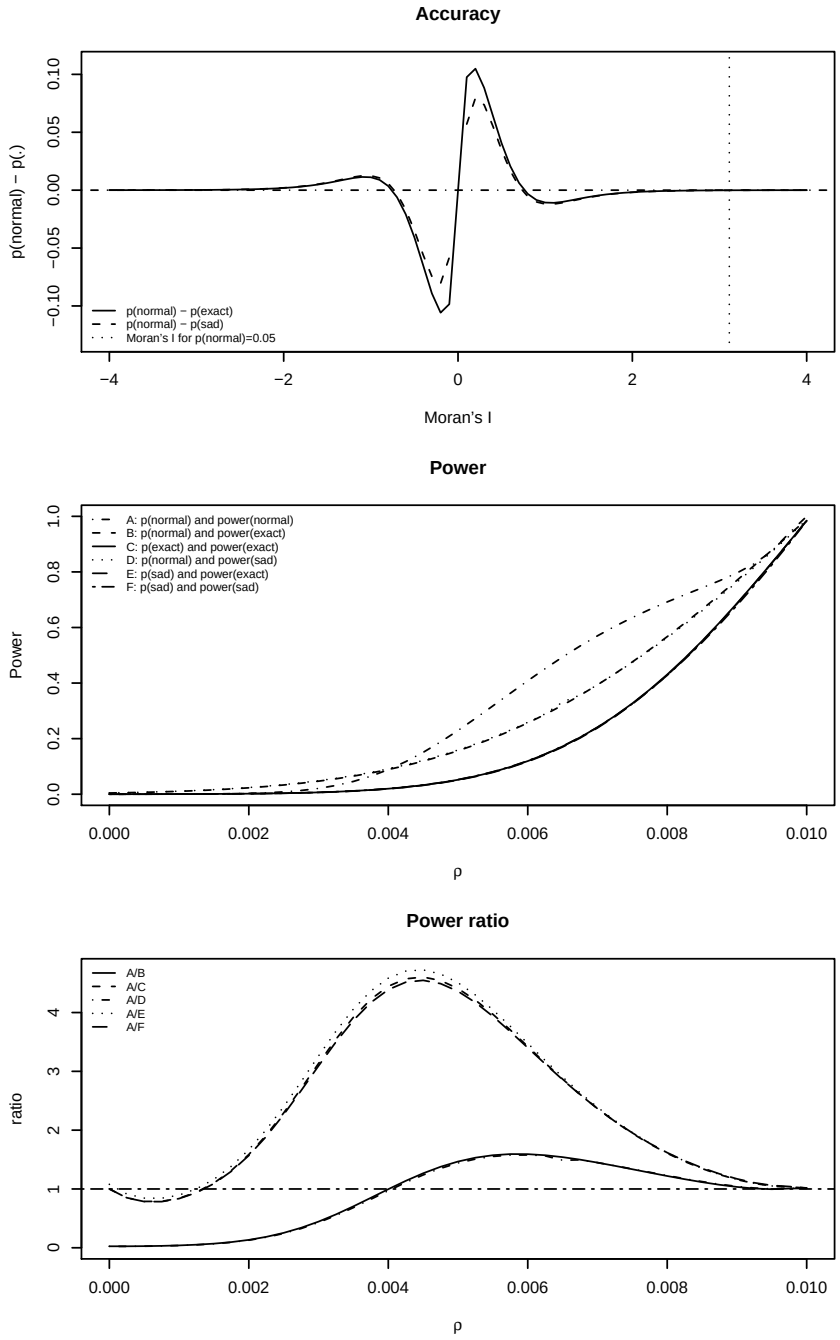

Figure 7: Atzesberg community, Upper Austria — upper panel: accuracy for different values of local Moran's I; centre panel: power for different values of ρ; lower panel: power ratios for different values of ρ.

moderate and larger values of ρ in the alternative hypothesis, although very small and very large values of ρ are not affected.

Once again, the power ratio lines form two groups, with the A/B and A/D lines lying close together — the Normal approximation test with power derived from the exact distribution and the Saddlepoint approximation respectively, and A/C, A/E, and A/F — the Saddlepoint approximation and Exact distribution based tests with power derived from the exact distribution (A/C, A/E) and the Saddlepoint approximation respectively (A/F). This result reinforces what we have learnt from analysing the B07 lattice, that on small lattices, be they global or star-shaped local lattices within maps of arbitrary size, the normal approximation may be misleading, and that alternatives will provide much better guidance.

5 Conclusions

In this paper, we have investigated the effect of using the common normal approximation for the distribution of global and local Moran's I on power calculations. Our substantive cases demonstrate that the induced errors can be severe and lead to be over- and rarely also underconfidence when testing for local clusters and hotspots.

Since the Saddlepoint approximation also fares well, with today's computing facilities there are no reasons left for using the normal approximation. The analyst can choose between the Saddlepoint approximation and the exact distribution, and as we have shown, these run at about the same speed because they use the same analytical pair of eigenvalues. For this purpose we have provided code for the exact distribution in the **Rspdep** package. Because both the exact distribution Saddlepoint approximation code now provide internal functions and intermediate objects, such as star-shaped weights structures and eigenvalues, power analysis has been made easier in practice, and we hope this contributes to making power analyses a regular feature of investigations of spatial patterns.

Note that the calculation of power function values may also be used in other contexts. Gumprecht et al. (2008) for instance employs them for defining a criterion to select sample points. They still use normal approximations and their method would greatly benefit from taking into account our findings.

Appendix

Full code can be found in Reder (2007). R-scripts for all the examples are contained in the file `bmr_scripts.zip` downloadable close to this technical report.

References

Ainsworth, L.M., Dean, C.B., 2008. Detection of local and global outliers in mapping studies. Environmetrics 19, 21–37.

Anselin, L., 1995. Local Indicators of Spatial Association - LISA Geographical Analysis 27, 93-115.

Bivand, R. and contributors, 2008. spdep: Spatial dependence: weighting schemes, statistics and models. R package version 0.4-15. http://cran.r-project.org/src/contrib/Descriptions/spdep.html

Boots, B. N., Royle, G. F., 1991. A Conjecture on the Maximum Value of the Principal Eigenvalue of a Planar Graph. Geographical Analysis 23, 54–66.

Cliff, A. D., Ord, J. K., 1973. Spatial autocorrelation. Pion, London.

Dhrymes, P. J., 1984. Mathematics for Econometrics. Springer, New York.

Goovaerts, P., Jacquez, G. M., 2005. Detection of temporal changes in the spatial distribution of cancer rates using local Moran's I and geostatistically simulated spatial neutral models. Journal of Geographic Systems 7(1) 137–159.

Gumprecht, D., Müller, W. G., Rodríguez-Díaz, J. M., 2008. Design for detecting spatial dependence. Geographical Analysis forthcoming.

Hepple, L. W., 1998. Exact testing for spatial autocorrelation among regression residuals. Environment and Planning A 30, 85–108.

Imhof, J. P., 1961. Computing the Distribution of Quadratic Forms in Normal Variables. Biometrika 48, 419–426.

Jacqmin-Gadda, H., Commenges, D., Nejjari C., Dartigues, J.-F., 1997. Tests of geographical correlation with adjustment for explanatory variables: an application to dyspnoea in the elderly. Statistics in Medicine 16, 1283–1297.

Leung, Y., Mei, C.-L., Zhang, W.-X., 2003. Statistical test for local patterns of spatial association. Environment and Planning A 35, 725–744, 2003.

Li, H., Calder, C. A., Cressie, N., 2007. Beyond Moran's I: Testing for spatial dependence based on the spatial autoregressive model. Geographical Analysis 39, 357–375.

Lloyd C. D., 2007. Local Models for Spatial Analysis. CRC Press, Boca Raton.

McLaughlin, C. C., Boscoe, F. P., 2007. Effects of randomization methods on statistical inference in disease cluster detection. Health & Place 13, 152–163.

Ord, J. K., Getis A., 2001. Testing for local spatial autocorrelation in the presence of global autocorrelation. Journal of Regional Science 41(3), 411–432.

R Development Core Team, 2007. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. http://www.r-project.org/

Reder, M., 2007. Global and local Moran's I's. An Improvement of the Geostatistical R-package spdep. unpublished Master-thesis at the Johannes-Kepler-University of Linz, Austria.

Schmoyer, R. L., 1994. Permutation tests for correlation in regression errors. Journal of the American Statistical Association 89(428), 1507–1516.

Tiefelsdorf, M., 1998 Some practical applications of Moran's I's exact conditional distribution. Papers in Regional Science 77(2), 101–129.

Tiefelsdorf, M., 2000. Modelling Spatial Processes, The Identification and Analysis of Spatial Relationships in Regression Residuals by Means of Moran's I. Springer, Berlin.

Tiefelsdorf, M., 2002. The Saddlepoint Approximation of Moran's I's and Local Moran's I_i's Reference Distributions and Their Numerical Evaluation. Geographical Analysis, 34, 187–206.

Tiefelsdorf, M., and Boots, B. N., 1995. The exact distribution of Moran's I. Environment and Planning A 27, 985–999.

Tiefelsdorf, M., and Boots, B. N., 1997. A Note on the Extremities of Local Moran's I and Their Impact on Global Moran's I. Geographical Analysis 29(3), 248–257.