Journal of the Statistical and Social Inquiry Society of Ireland
Vol. XXXVIII

EXPLORING THE ECONOMIC GEOGRAPHY OF IRELAND

Edgar Morgenroth[1]

Economic and Social Research Institute

(read before the Society, December 11th, 2008)

***Abstract*:** Only a few research papers have analysed the spatial distribution of economic activity in Ireland. There are a number of reasons for this, not least the fact that comprehensive data on the location of economic activity by sector across all sectors has not been available at the highly disaggregated spatial level. This paper firstly establishes the geographic distribution of employment at the 2 digit NACE level, using a novel approach that utilises a special tabulation from the CSO 2006 Census of Population Place of Work Anonymised Records (POWCAR). It then analyses the spatial patterns of this distribution using maps and more formal methods such measures of spatial concentration and tests for spatial autocorrelation. The paper considers the locational preferences of individual sectors, the degree to which specific sectors agglomerate and co-agglomerate, and thus will uncover urbanisation effects and differences across urban and rural areas regarding economic activity.

Keywords: Economic geography, employment distribution,
JEL Classifications: R12, R14, R30.

1. INTRODUCTION

In modern economies economic activity is not spread evenly across space. This is also holds in Ireland, yet the spatial distribution of economic activity in Ireland has received relatively little attention by researchers. The few papers that have been published have tended to focus on manufacturing alone or have other drawbacks.

For example Strobl (2004) measured the degree of localisation of economic activity. In particular he analysed the spatial distribution of broad sectors over the period 1926 to 1996 using Census and Forfas Employment Survey data. However, the Census data referred to the residential location of individuals employed in different sectors. Given the substantial degree of commuting across regions, counties and particularly Electoral Districts, this analysis may be subject to substantial bias.

Gleeson, Ruane and Sutherland (2006) used data from the Census of Industrial Production to analyse the sectoral specialisation and spatial dispersion of manufacturing activity. Their analysis focused particularly at the distinction between indigenous and multinational enterprises (MNEs) and at changes between the period 1985-1993 (pre Celtic Tiger) and 1993 to 2002 (Celtic Tiger). They found increasing specialisation and spatial concentration of MNE employment but less specialisation and more dispersal for indigenous employment.

Morgenroth (2008a) applied Krugman's measure of relative spatial specialisation to employment data at the regional authority (NUTS 3) for manufacturing sectors drawn from the Forfas employment survey. Overall, his results identify a decline in specialisation but this is less pronounced in peripheral regions. His results also identify a negative relationship between specialisation and growth.

The primary reason for little research being done in this area is that the required detailed data of the location of economic activity has not been available. It is also notable that this area of research has not been popular among geographers (there are some exceptions, e.g. Breathnach, 2000). All of the previous papers that have considered the economic geography at a wider sectoral level have been written by economists. In addition, there has been some research that has focused only on specific sectors (e.g. van Egeraat and Jacobson 2005, 2006).

[1] The author would like to thank the CSO, and particularly Gerry Walker, for making available the data used in the analysis.

This paper aims to fill the research gap by calculating the employment (jobs) at the Electoral Division (ED) level, disaggregated by two digit NACE sector, including a detailed disaggregation of the services sector. This is achieved using a novel approach by utilising a special tabulation from the CSO 2006 Census of Population Place of Work Anonymised Records (POWCAR). Having established the spatial distribution of employment by sector, the paper analyses the spatial patterns of this distribution using a number of spatial statistical methods such as tests for spatial autocorrelation. This analysis uncovers the locational preferences of individual sectors, the degree to which specific sectors agglomerate and co-agglomerate, and thus shows the degree of urbanisation effects and differences across urban and rural areas regarding economic activity.

The interest by economists in the spatial distribution of economic activity derives from the growth of what has become known as the "New Economic Geography" literature, which was initiated by Paul Krugman in a series of influential papers (e.g. Krugman, 1991). This literature highlights the importance of agglomeration economies/diseconomies in driving the spatial pattern of economic activity. There has been an ongoing interest in testing the predictions of this literature, which was at least initially overly theoretical, in empirical studies.

There has also been substantial interest in the drivers of the differential growth rates across regions. In particular, some authors have highlighted the role of specialisation in particular sectors as determining regional growth rates (see Paci and Pigliaru, 1999). Of course the argument that regions specialised in low growth sectors necessarily will have low growth rates is tautological. Nevertheless, a high level of specialisation is likely to increase the volatility of growth rates at the regional level as more specialised regions are more susceptible to shocks (lack of diversification).

Finally, in order to analyse the impact of economic structural change on local labour markets or commuting behaviour it is necessary to establish the economic geography of a country. Such an analysis is carried out in a separate paper (Morgenroth, 2008b).

This paper is organised as follows. Section 2 outlines the data sources and the construction of the job numbers by Electoral District. Section 3 aims at identifying the broad spatial patterns using maps. Section 4 explores the data more formally using a range of tools. The final section summarises the findings and offers some policy implications.

2. DATA

The analysis conducted here draws almost entirely on one data set. This is a special tabulation from the CSO 2006 Census of Population Place of Work Anonymised Records (POWCAR).

The Census has included a question on the address of the place of work of respondents at least since 1986. However, this information was not used for Censuses before the 2002. Following the 2002 Census, the CSO geocoded the place of work of respondents who were enumerated in a private household, were 15 years old or over, were enumerated at home and indicated that their Present Principal Status was working for payment or profit. In 2002, the place of work was geocoded for 15% of respondents and the data was made available as the Place of Work Sample of Anonymised Records (POWSAR). For the 2006 Census, the CSO attempted to geocode the place of work of 100% of records, with the resulting data having been made available as Place of Work Census Anonymised Records (POWCAR)

As the name suggests, POWCAR is a microdataset which contains a range of variables at the individual level, including the sector in which individuals work. Other variables include basic demographic data such as gender and household composition, details about the nature of the commuting behaviour such as mode, time of leaving home, distance and time taken, and other basic background variables such as education. POWCAR records the ED in which the individual is resident (which is recorded by the Census enumerator) along with the ED in which the place of work is located. The latter is achieved by geocoding the stated place of work. This is only done for individuals who have not indicated that they work from home, have no fixed workplace (mobile worker) or have not filled in the address of the place of employment. For those individuals for which a place of employment was recorded it was matched against addresses on the An Post GeoDirectory. Where an exact match could not be found a near match was recorded.

While the POWCAR file, which is available to researchers under certain conditions, is extremely useful for a range of analyses, for present purpose it proved not to be ideal as the sectoral breakdown available in the file was too coarse (e.g. only 7 broad sectors are identified). On the other hand, POWCAR also contains a lot of individual level detail which is not needed for the present analysis. Consequently, a special tabulation was requested from the CSO which omitted all the individual level detail, but added additional detail on the sector.

The special tabulation contains counts of persons at work separately distinguished by electoral district (ED) of residence and ED of place of work disaggregated by 2 digit NACE sector.[2] The two series, persons at work by residence and by place of work are not linked, neither are other micro-variables included in the special tabulation. This preserves the anonymity of respondents and given the purpose of this paper, has no drawbacks for the analysis presented here.

Table 1 shows the number of workers by type place of work coding and the proportion of each category for which a NACE code was not identified. Overall, the cross tabulation provides by the type data for 1,834,472 individuals. Of those 6% work at home, 11% have no fixed place of work, for 75% the address of employment could be matched and for a further 8% the address was either not given or could not be matched. A small number (0.5%) work abroad of which the majority work in Northern Ireland. The Census identified the total numbers of employed persons as 1,930,042, which implies that POWCAR does not contain records for almost 100,000 workers. This is explained by the fact that the data only covers those in private households who enumerated at home. For most categories identified in the table the proportion for which a NACE code is available is very high. However, for those for which no address was available the proportion missing a NACE code is two thirds.

Table 1 Breakdown of POWCAR data

	Total	Missing NACE code
Place of Work stated	1,372,554 (74.8%)	8,694 (0.6%)
Home	107,202 (5.8%)	3,293 (3.1%)
Abroad	8,295 (0.5%)	325 (3.9%)
Mobile	208,548 (11.4%)	9,576 (4.6%)
Blank	137,873 (7.5%)	9,3203 (67.6%)
Total	1,834,472	115,091 (6.3%)

Source: Own calculations using POWCAR Special Tabulation

It is important to verify that the data concords with other data sources such as the Quarterly National Household Survey (QNHS), which provides detail of the total employment by broad sectors. Table 2 shows a comparison of the QNHS for quarter 2 of 2006 with the total derived from the POWCAR. The two data sets correspond well, with a correlation coefficient of 0.98. The primary difference is the proportion of the POWCAR based employment numbers which could not be attributed to a sector so that all but one sector shares is below that reported for the QNHS.

Table 2 Comparison between the POWCAR based and QNHS sectoral employment numbers

	QNHS	Share	POWCAR	Share
A-B Agriculture, Forestry and Fishing	114.5	5.7%	87.3	4.8%
C-E Other Production Industries	288.5	14.3%	250.4	13.7%
F Construction	262.7	13.0%	204.9	11.2%
G Wholesale and Retail Trade	284.4	14.1%	247.8	13.6%
H Hotels and Restaurants	116.3	5.8%	94.8	5.2%
I Transport, Storage and Communication	120.7	6.0%	101.1	5.5%
J-K Financial and Other Business Services	267.3	13.3%	251.5	13.8%
L Public Administration and Defence	105.1	5.2%	94.9	5.2%
M Education	135.6	6.7%	121.2	6.6%
N Health	201.2	10.0%	181.5	9.9%
O-Q Other Services	120.6	6.0%	75.9	4.2%
Not Classified			114.8	6.3%
All Economic Sectors	2017	100%	1826.2	100%

Note: in addition to the 1,826,177 (1826.2) employed persons working in the Republic of Ireland, the POWCAR identifies a further 8295 workers who work outside of the Republic of Ireland (largely in Northern Ireland).

[2] Given that the data is available at the ED level, the reference spatial unit in this paper unless explicitly highlighted is that of the ED.

In order to identify the total number of jobs located in any ED, the number of workers resident in that ED who reported that they were working from home is added to those number of jobs identified through the geocoding of the stated place of work. This leaves those who stated as having no fixed place of work and those who did not give an address for the place of work not counted into the number of jobs per ED. On closer examination this undercounts the number of jobs, particularly for the construction sector as 54% of those without a fixed place of work stated that they were working in the construction sector. Two options of attributing these mobile workers are possible. Their workplace could be attributed to their place of residence, or they could be attributed according to the distribution of the jobs identified precisely. Here we opt for the first solution, but the results of the analysis conducted below do not appear to be sensitive to the assumptions. Finally, so that the data adds to the total identified in the census the difference between the census and that accounted for by those with a place of work residence, home workers and mobile workers are attributed according to shares in employment where the NACE code is not available into that group. As the data for employment where the NACE code is not available is not used in the analysis below this attribution is of no consequence for the detailed analysis but obviously impacts on the total number of jobs in each ED.

Overall, the data covers 58 NACE sectors, but these are aggregated into 30 sectors. For example Agriculture and Forestry are identified separately. However, since they are usually aggregated into on sector we follow this convention here.

3. MAPPING SPATIAL DISTRIBUTIONS

The most natural way to identify the spatial distribution of sectoral employment is to map the data. However, given the fact that EDs differ substantially in size it is necessary to scale the data appropriately, by converting the absolute job numbers into a density of jobs per square kilometre. This also has the advantage that this also results in a further “anonymisation”.

In the first instance, it is quite instructive to compare the density of employment with that of the population which are shown in Map 1 and Map 2 respectively. In both maps the data intervals are identical. It is immediately obvious that employment is considerably more concentrated spatially than the population. Indeed, the population has been dispersing over recent years. Also obvious in the maps is the concentration of both employment and the population in and around the major urban centres. In relation to the population the maps clearly shows the upland areas that are essentially unpopulated, the rural areas (10 to 50 persons per km^2) and the small urban and peri-urban areas. With respect to the latter, it is noticeable that they are adjacent to the major urban areas.

Of course not all sectors have the same locational preferences and hence it is useful to consider the distribution of employment in individual sectors. Such an analysis can uncover clustering or urbanisation driven agglomeration along with the co-location of individual sectors. Given the way our data is constructed this is a straightforward task. In order to conserve space further maps are produced for a few representative sectors only. These are Agriculture and Forestry, Food and Drink, Chemicals and Chemical Products, Construction, Financial Services and Education. These sectors cover primary production modern manufacturing, traditional manufacturing market services and public services. The maps for these sectors can be found at the end of this paper. Since the sectors differ in total size it is difficult to show the data using the same intervals for each sector. However, we use a the dot density (choropleth) map with a dot representing a single job.

Comparing the maps for these sectors (Map 3 to Map 8: located at the end of the paper) shows that there are marked differences between the sectors with respect to the spatial distribution of employment. Among this sample Agriculture and Construction are the most dispersed sectors while the other sectors are more concentrated, indicating a strong relationship with urban areas. Chemicals and Financial services appear to be most concentred. Food and Drink appear to be particularly strong in Munster, Leinster and Monaghan.

Map 1 Population Density (persons per km^2), 2006

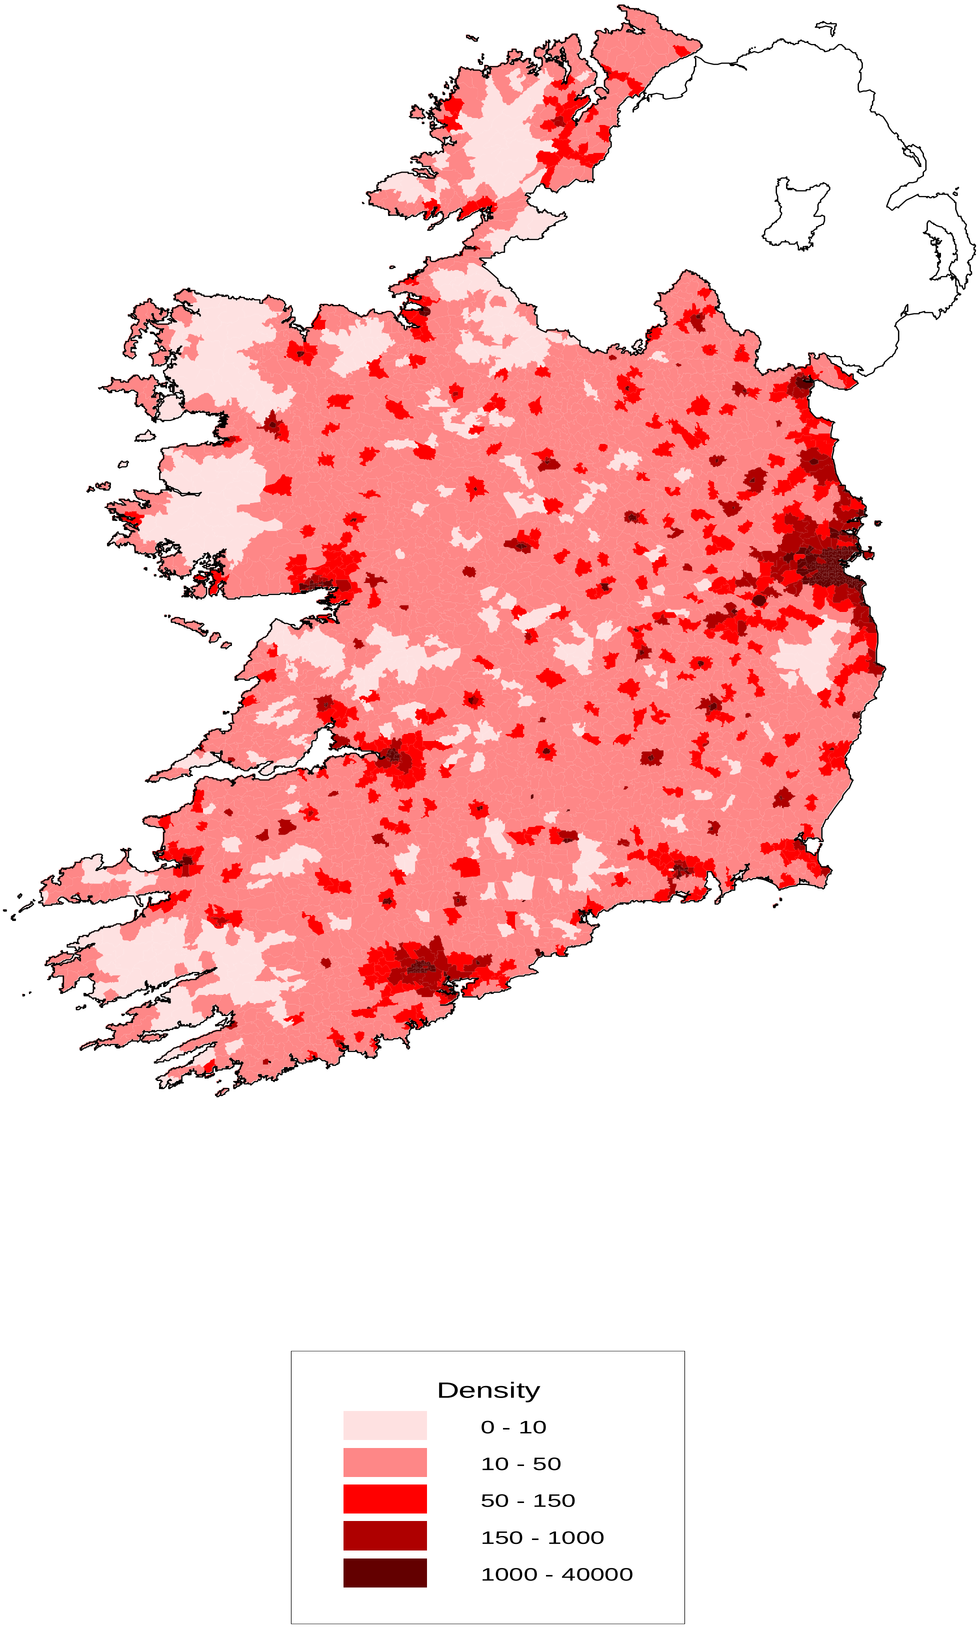

Source: Own calculation using CSO Census 2006, SAPS.

Map 2 Job Density (jobs per km^2), 2006

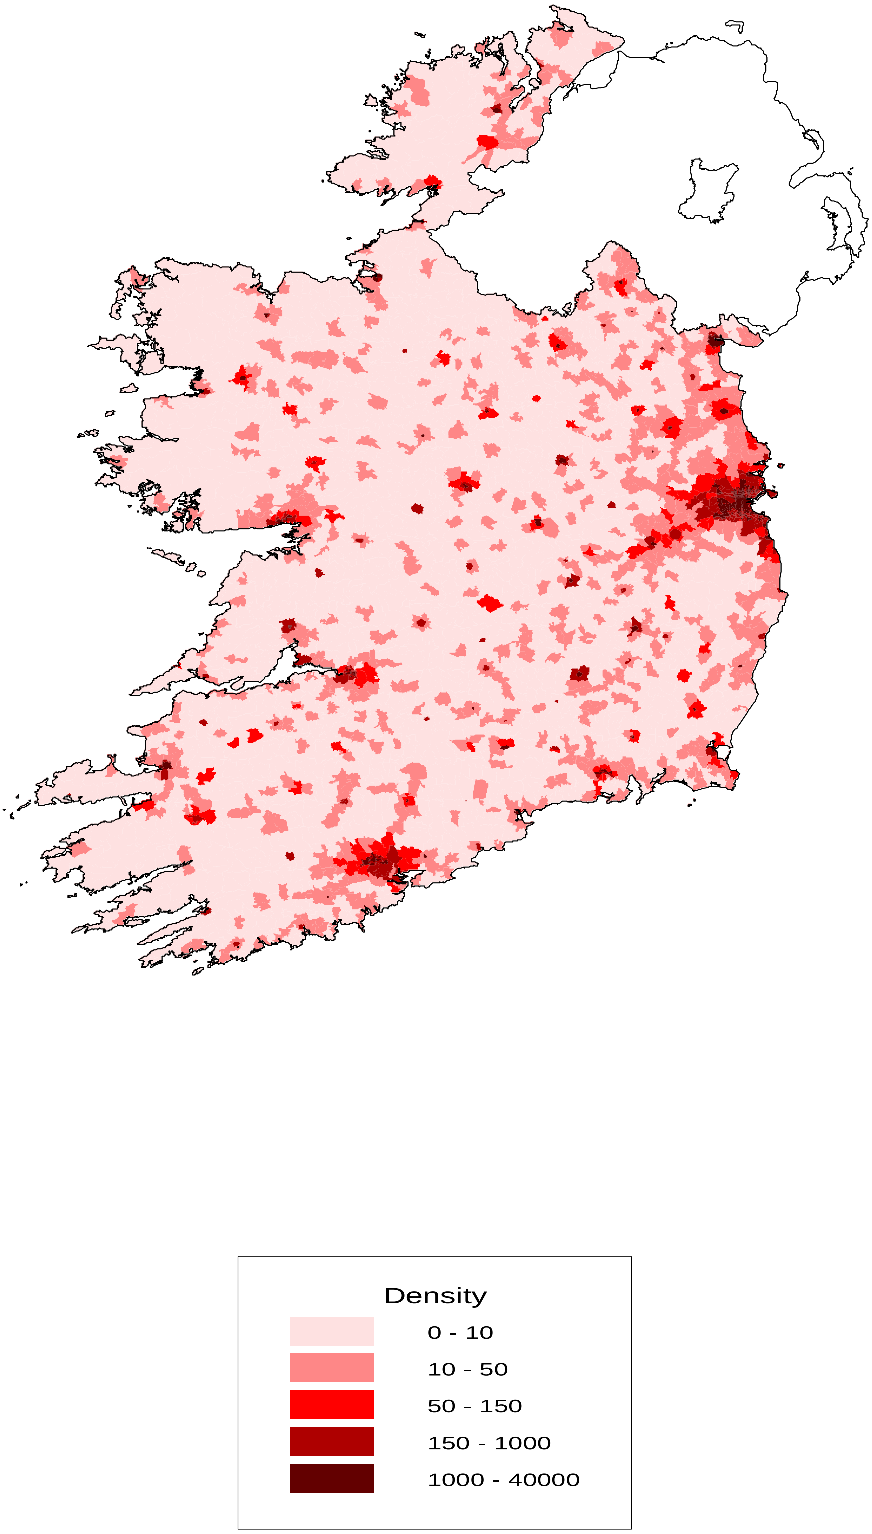

Source: Own calculations using POWCAR Special Tabulation

4. FORMAL TESTING OF SPATIAL DISTRIBUTIONS

While maps are powerful tools for visualising spatial patterns, they are not without limitations. Firstly, the information conveyed in a map is not independent of class sizes/dot size etc. and are thus subject to the criticism that they are social constructs (Crampton, 2001). Secondly, maps are subject to what is known as the modifiable area unit problem (MAUP). Thus, as Gehlke and Biehl (1934) outline, heterogeneity across space implies that the results from any analysis are likely to depend on the nature and degree of aggregation across spatial units (similar aggregation problems exist in economics e.g. sectors). Finally, maps have a limited use in identifying and quantifying the underlying processes.

Consequently, more formal methods of describing the spatial distribution of employment need to be utilised. A number of alternative methods will be used here. Firstly, we calculate spatial Herfindahl indices, which have previously been applied in an Irish context by, Morgenroth (2008a) and Gleeson, Ruane and Sutherland (2006).[3] These measures provide an indication of the spatial concentration of employment in different sectors.[4]

Secondly, measures of spatial autocorrelation are used to identify the degree to which employment densities in individual sectors are correlated across spatial units. Given that the correlation across spatial units is multidimensional, conventional correlation coefficients cannot be applied. Rather special spatial autocorrelation tests are used here.

Before we can define the measure of spatial concentration it is useful to define the key variables used in its construction:

E_{ij} - is employment in sector *i* in ED *j*, where $i = 1,2\ldots I$ and $j = 1,2,\ldots R$

$E_i = \sum_j E_{ij}$ is total employment in industry *i*. $E_j = \sum_i E_{ij}$ is total employment in ED j. We can define the share of employment in sector *i*, in region *j* in total sectoral employment is given as:

$$c_{ij} = \frac{E_{ij}}{E_i} = \frac{E_{ij}}{\sum_j E_{ij}}$$

Using this we can define the Herfindahl index of absolute specialisation as:

$$H_i^C = \sum_j (c_{ij})^2 \quad \text{which takes values} \quad \frac{1}{R} \le H_i^C \ge 1$$

In words, the Herfindahl index of absolute concentration is defined as the sum of the squared shares of the regional sectoral employment for each sector. Using the same approach it is also possible to calculate a Herfindahl index of specialisation of each ED which is constructed by calculating the sum of squared employment shares by sector for each ED.

In order to derive the measures of spatial correlation it is necessary to define the structure of the spatial relationships between EDs. This is achieved through the use of a spatial weights or connectivity matrix, W, consisting of individual elements w_{ij} and where the diagonal elements are equal to zero. An important issue is the choice of the weights, w_{ij}. One of the most widely used specification of these spatial weights is based on the concept of connectivity which is measured as a binary variable which is equal to one if EDs i and j have a common border and zero if they do not have a common border.[5] This implies that such a specification assumes that only neighbouring EDs are taken into account when measuring the correlation across spatial units. Another widely used specification utilises the distance or inverse distance between two EDs, which implies a distance decay of the relationship (see Ord, 1975; Cliff and Ord, 1981). This latter approach of the spatial weights matrix has the advantage of satisfying Tobler's first law of geography that "everything is related to everything else, but near things are more related than distant things" (see Tobler, 1970). The drawback of the latter approach is that for the number of spatial units used here (3,441) results in a very large matrix which is difficult to handle even with significant computing resources. In the present work it would result in a 3440 by 3440 matrix i.e. a matrix with 11,833,600 elements! Consequently, the chosen weights matrix for the analysis

[3] Strobl (2004) used a slightly different measure namely a Gini coefficient based location quotients to measure the degree of localisation.

[4] A range of measures of spatial concentration have been proposed in the literature (Bickenbach and Bode, 2008).

[5] Moran (1948) and Geary (1954) first proposed binary contiguity between spatial units in their pioneering papers on measures of spatial dependence.

conducted here is of the binary contiguity type. In order to allow for correlation beyond the immediate neighbours secondary contiguity is allowed.[6]

Two measures of spatial correlation are applied here namely, Moran's I (see Moran, 1948), and Geary's C (see Geary, 1954). Formally, Moran's I is given as:

$$I = \frac{N}{S_0} \frac{\sum_i \sum_j w_{ij} (x_i - \mu)(x_j - \mu)}{\sum_i (x_i - \mu)^2}$$

Where N is the number of observations, w_{ij} is the element in the spatial weights matrix corresponding to the pair of observations ij, x_i and x_j respectively are observations at locations i and j and S_0 is a scaling constant equal to the sum of all weights.[7] The mean of the observations x is denoted by μ. This coefficient, while similar to a correlation coefficient, is not centred around 1. Rather the expected value of I is negative and depends on the sample size with that expected value tending towards zero as the sample size increases. A Moran's I less than expected implies negative spatial autocorrelation while one larger than expected implies positive spatial autocorrelation. The second measure, Geary's C is given as:

$$C = \frac{N-1}{2S_0} \frac{\sum_i \sum_j w_{ij} (x_i - x_j)^2}{\sum_i (x_i - \mu)^2}$$

In contrast to the more widely used Moran's I, Geary's C has an expected value of 1 (is centred around 1). In contrast to the standard correlation coefficient, a value greater than one indicates negative spatial correlation while a value lower than one indicates positive spatial correlation. Moran's I has become the most widely used measure of spatial autocorrelation since it is less affected by deviations of the sample data from the standard normal distribution (see Cliff and Ord, 1981).

Table 3 shows the statistics for each measure along with the number of EDs which have any jobs in the respective sector. This latter indicator is important in interpreting the results, particularly for the Herfindahl index since a large number of ED without employment in the sector implies that the remaining EDs have a relatively large share of the total sectoral employment and hence a higher Herfindahl index. Of course, the number of EDs with employment is related to the overall size of the sector. Sectors with a large number of jobs should be represented in more EDs. This is what has been referred to as the dartboard effect – if one has lots of darts they will land in a larger number of fields than if one only has a few. More formally, the relationship between sectoral size and numbers of EDs can be captured through a correlation coefficient, which for our data turns out to be 0.79, which indicates a strong positive relationship.

Turning to the indicators, the Herfindahl index which measures the degree to which employment in each sector is concentrated across EDs is found to be large spread, ranging from 0 for agriculture and forestry to 0.28 for fuels. However, the latter sector, due to the small number of ED which contain jobs of this sector and the small overall size of the sector, is a significant outlier. If one ignores this sector the variance is much reduced but as Figure 1 shows, there are still substantial differences between sectors. The manufacture of transport equipment (NACE 34-35) is found to be the most concentrated sector followed by Electrical and Optical Equipment (NACE 30-33), Financial Services (NACE 65-67) and Electricity, Gas and Water Supply (NACE 40-41). The least concentrated sectors are Agriculture and Forestry (NACE 1-2), Construction (NACE 45) and Sale and repair of Motor Vehicles (NACE 50).

The degree of specialisation of EDs is calculated using a Herfindahl index defined over EDs as outlined above. This yields an index for each ED, which is best displayed in a map. Map 9 shows the results of this calculation. The Herfindahl index of ED specialisation ranges from 0.066 to 0.76. In other words the most specialised ED is more than 10 times more specialised than the most diversified ED. The map shows an interesting spatial pattern of specialisation. EDs surrounding the major urban areas are the least specialised while some urban EDs are very specialised and many rural EDs have either a high or medium level of specialisation. Clearly visible in the map are EDs with a known level of high specialisation. For example in North Dublin, Airport ED is very highly specialised (in Transport Storage and Communications). Of course the types of industry that dominate in the more specialised EDs vary significantly. In rural areas more traditional activities predominate, while in urban areas services or more modern manufacturing dominates.

---

[6] The spatial weights allow secondary contiguity and are of the 'queen' type.

[7] With a row standardise spatial weights matrix (achieved by dividing each element of the spatial weights matrix by its row total) the scaling constant equals the number of observations so that Moran's I simplifies slightly.

Table 3 Formal Measures of Spatial Concentration and Spatial Correlation

NACE Code	Sector Description	EDs with employment	Herfindahl	Moran's I	Geary's C
1-2	Agriculture and Forestry	3,354	0.000	0.158***	0.76***
5	Fishing	469	0.014	0.041***	0.93**
10-13	Mining and Quarrying	1,221	0.014	0.025***	0.87*
15_16	Manufacture of Food and Drink	2,016	0.006	0.031***	0.96
17-19	Manufacture of Textiles and Leather	1,060	0.008	0.185***	0.89***
20	Manufacture Wood and Wood Products	1,138	0.006	0.092***	0.99
21	Manufacture of Paper and Paper Products	412	0.021	0.047***	0.99
22	Publishing	1,090	0.008	0.172***	0.88**
23	Manufacture of Fuels	74	0.280	0.029***	0.87*
24	Manufacture Chemicals and Chemical Products	1,187	0.021	0.026**	1.00
25	Manufacture of Rubber and Plastic	731	0.011	0.025***	0.86**
26	Manufacture of Non-metallic Minerals	1,908	0.006	0.013***	0.94
27-28	Manufacture of Basic Metals and Fabricated Metal Products	2,318	0.004	0.196***	0.78***
29	Manufacture of Machinery and Equipment	1,468	0.010	0.039***	0.98
30-33	Manufacture of Electrical and Optical Equipment	1,327	0.036	0.066***	0.90
34-35	Manufacture of Transport Equipment	656	0.049	0.004*	1.02
36-37	Manufacturing not elsewhere classified	1,969	0.004	0.175***	0.79**
40-41	Electricity, Gas and Water supply	1,402	0.021	0.055***	0.95
45	Construction	3,435	0.001	0.564***	0.45***
50	Sale and Maintenance of Motor Vehicles	2,643	0.002	0.257***	0.76***
51	Wholesale	2,823	0.004	0.250***	0.76***
52	Retail	2,776	0.005	0.100***	0.89**
55	Hotels and Restaurants	2,678	0.003	0.200***	0.79**
60-64	Transport, Storage and Communications	3,163	0.009	0.265***	0.74**
65-67	Financial Intermediation	1,592	0.023	0.235***	0.83**
70-74	Real Estate, Renting and Business Activities	3,062	0.007	0.345***	0.72***
75	Public Administration and Defence	2,567	0.008	0.232***	0.83***
80	Education	3,037	0.003	0.295***	0.68***
85	Health and Social Work	3,143	0.005	0.163***	0.82***
90-99	Other Community, Social and Personal Services	2,933	0.003	0.325***	0.65***

Note: *, **, *** indicate significance at the 90%, 95% and 99% level respectively.

The measures of spatial autocorrelation, which identify the degree to which the employment density for a particular sector is correlated across spatial units, shows positive spatial autocorrelation. That is, EDs with a high density are typically surrounded with EDs which have similar density in that sector. Both measures are highly correlated with a correlation coefficient of -0.97. The negative correlation derives from the fact that for Geary's C, a statistic that is smaller than one indicates positive spatial autocorrelation while for Moran's I a larger value indicates positive spatial autocorrelation. Interestingly, the degree of statistical significance of the measures is significantly lower for Geary's C.

For Moran's I, all statistics indicate positive spatial correlation and all but one are statistically significant at the 95% level. On the other hand one of the Geary statistics indicates a negative spatial correlation and only 19 of the 30 statistics are significant at the 95% level. However, there is a high degree of concordance regarding ranks between both statistics. The most autocorrelated sector is Construction, while the least autocorrelated sector is the Manufacture of Transport Equipment. This suggests an interesting relationship between the measure of spatial concentration and that of spatial autorrelation, in that the least concentrated sectors appear to be most spatially correlated. This relationship is confirmed by correlation coefficients but is not as strong as one would expect.

Figure 1 Herfindahl Index of Spatial Concentration

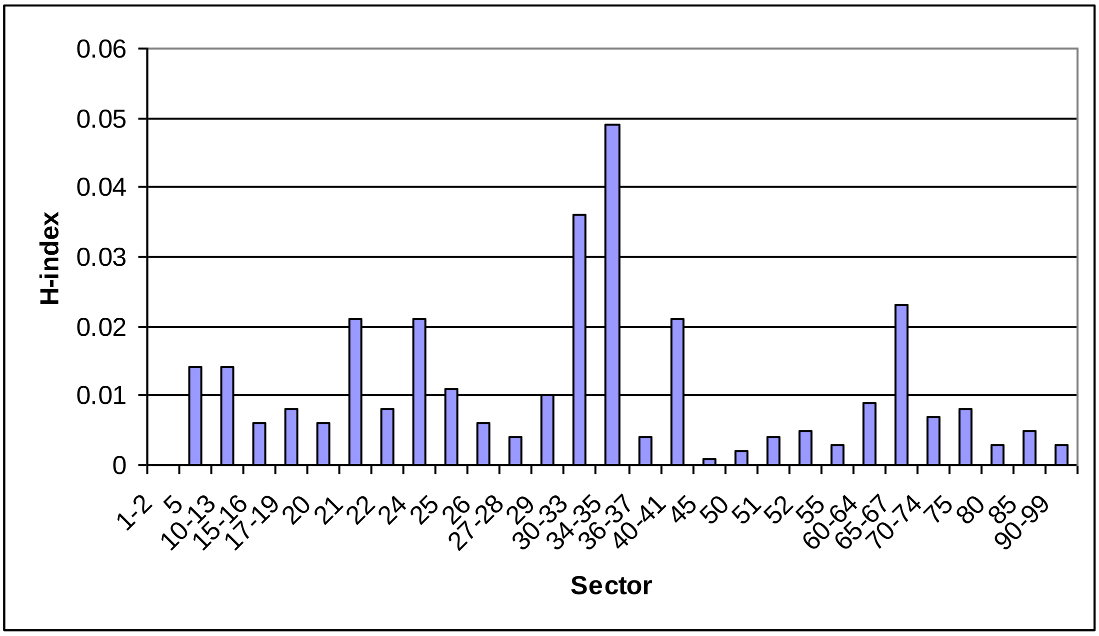

Source: Own calculations

Map 9 Herfindahl Index of ED specialisation

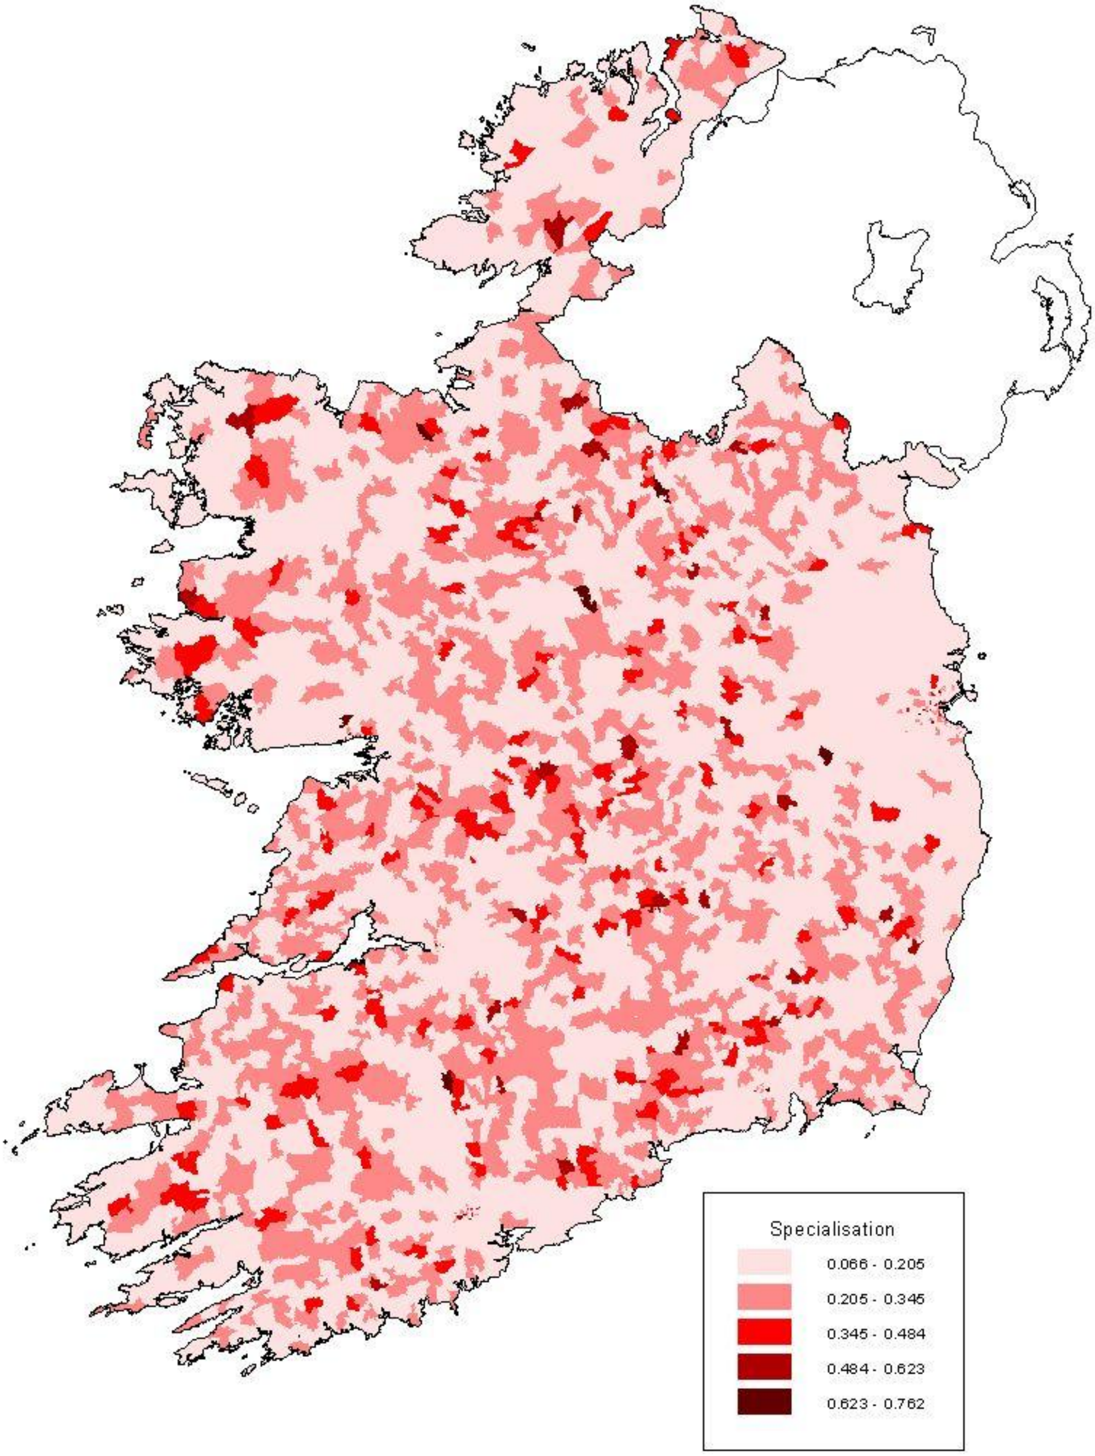

Source: Own calculations using POWCAR Special Tabulation

Measures such as those used above, while useful in identifying the overall spatial aggregation and concentration, are not able to account for spatial non-stationarity (heterogeneity), nor are they able to identify any local spatial clustering.

Anselin (1995) suggests a measure of local indicators of spatial association (LISA) which decomposes the Moran statistic into a local Moran statistic to identify the degree of spatial clustering and the contribution of each spatial unit towards the global Moran statistic. Formally, the local Moran statistic is given as:

$$I_i = (x_i - \mu)\sum_j w_{ij}(x_j - \mu)$$

The subscript indicates that this statistic is calculated for each individual spatial unit. These statistics are best displayed through mapping. Again for lack of space the results for only a few representative sectors are shown (Maps 10 to 15). These maps show that the global Moran's I statistics are largely due to spatial correlation across similar EDs with low densities, with the exception of Dublin where for many sectors there is a high correlation across high density EDs. The latter points to urbanisation economies while the latter identifies areas with very low levels of economic activity.

In the case of Agriculture and Forestry a cluster of high density EDs can be identified particularly in north Dublin, which is likely to be explained by a concentration of market gardening, which has a high labour intensity and hence a larger number of jobs. For Food and Drink, an interesting low-high cluster can be identified around Carrickmacross in Co. Monaghan, which indicates a high density cluster in the town with neighbouring EDs lacking employment in that sector. For Chemicals and Chemical Products a number of high-high clusters can be identified, especially in Dublin and Cork. For Construction the primary high-high cluster is in Dublin reflecting the density of larger construction projects. For the Financial Services sector the principle high-high cluster is in Dublin, corresponding to the IFSC and reaching into Dublin 2 and Dublin 4. Finally, for Education, the high-high clusters are found in the cities with universities, which of course also have a large number of schools reflecting their population, and particularly Dublin.

The basic mapping and the LISA analysis suggest a significant difference between urban and rural areas with respect to the type of economic activity present. In order to further analyse this data we apply some basic regression analysis which also helps in identifying the degree of concentration. Rather than using the 'raw' densities it is useful to standardise the densities for each sector by dividing them by the largest density. It is then possible to sort the EDs by their standardised density and compare this across sectors. Graphing this data yields employment density gradients which are shown in Figure 2, where the y-axis scale is logarithmic. The flattest curve, indicating a relatively even distribution of densities, is that for Agriculture and Forestry, while the steepest curve is that for fuels, which of course is a sector that is present only in a few EDs reflected in the fact that the curve is very short. Other relatively concentrated sectors are Manufacture of Transport Equipment, Manufacture of Paper and Paper Products, Electricity, Gas and Water Supply and Financial Services. The most dispersed sectors include Construction, Sale and Repair of Motor Vehicles, Manufacture of Basic Metals and Fabricated Metal Products and Education. These results confirm those of the analysis above.

Figure 2 Employment Density Gradients

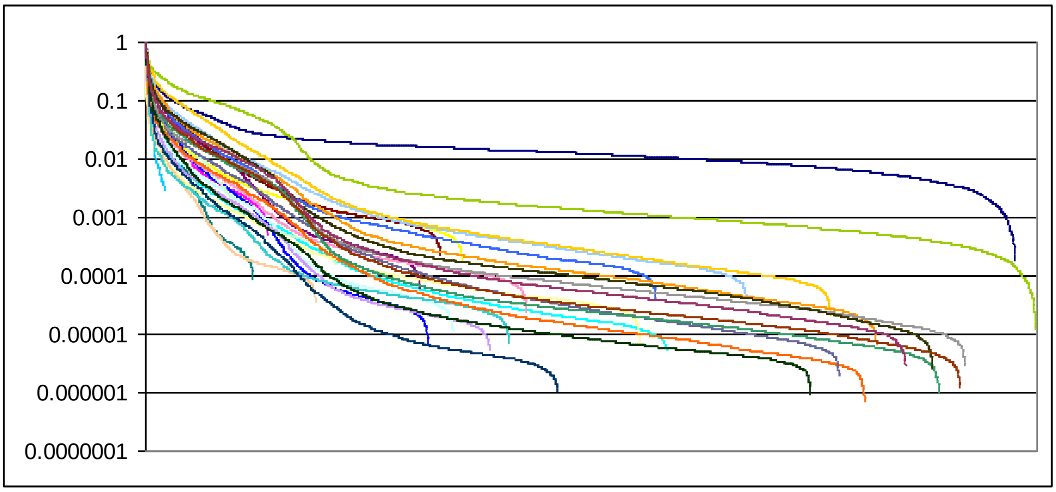

Source: Own calculations. Y-axis scale is logarithmic.

In order to estimate the slope of these employment density gradients it is straightforward to apply a simple model relating the density of jobs in each sector to the rank in the density distribution. More formally, taking logs this is given as:

$$\log(Density) = \log A - \alpha \log Rank$$

This relationship can be readily extended by letting the density also depend on whether the ED is urban or not, by adding a dummy variable for this, so that the relationship becomes:[8]

$$\log(Density) = \log A - \alpha \log Rank + \beta Urban$$

Given the large number of sectors it is more instructive to display the results in graphical form, where the point estimates for the parameters are shown as points and the confidence interval of two standard deviations is delineated by a high and low horizontal bar. In Figure 3, the slope parameter from the regression of density on rank is shown. In Figure 4, the corresponding parameter from the regression including the urban dummy is shown and finally, in Figure 5, the parameters for the urban dummy are shown.

Figure 3 clearly shows that the slopes for many sectors are statistically different from each other. In particular those of the services sectors are uniformly steeper, while those of the primary sectors are flatter. A mixed picture emerges for the manufacturing sector with Publishing and Electrical and Optical Equipment having steep density gradients, while Wood and Wood Products and Fuels have relatively flat slopes. Once the urban dummy is added to the regression model, the slopes flatten in all cases and the variation between sectors reduces significantly. Nevertheless, there are still a significant number of slopes which are statistically significantly different. The fact that the variation in the slopes reduces with the addition of an urban dummy highlights the importance of urban location in a number of sectors. These are identified in Figure 5, which shows the size of the coefficient for the urban dummy. Again the services sectors are noticeable for having a large parameter as does Publishing which is naturally found in urban areas given demand linkages, and Electrical and Optical Equipment. Overall, the more traditional manufacturing sectors and the primary sectors are less affected by urbanisation.

Figure 3 Slopes from Regressing log of Density on log Rank

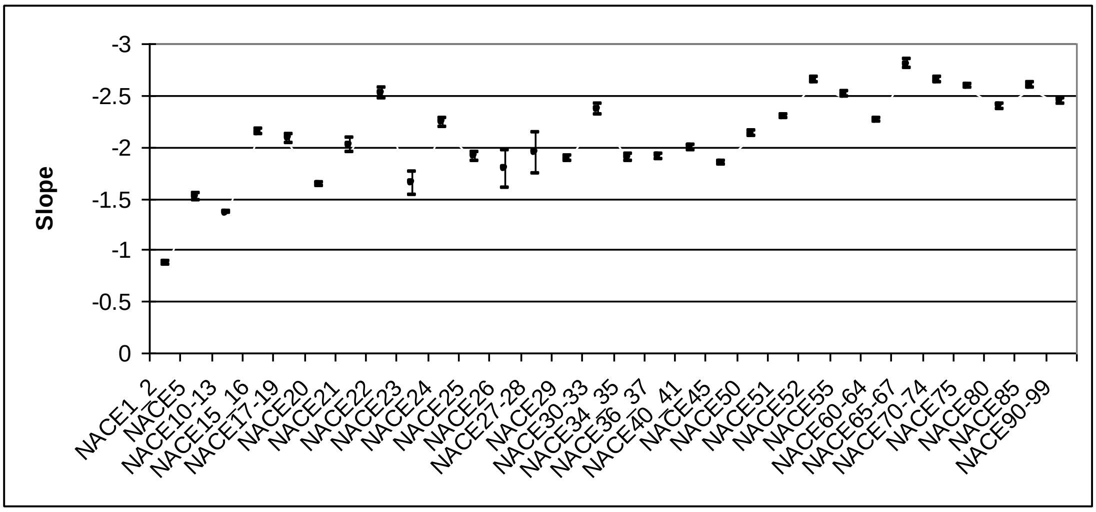

Source: Own calculations

[8] Urban EDs are those with a population density in excess of 150 persons per square kilometre.

Figure 4 Slopes from Regressing log of Density on log Rank & Urban

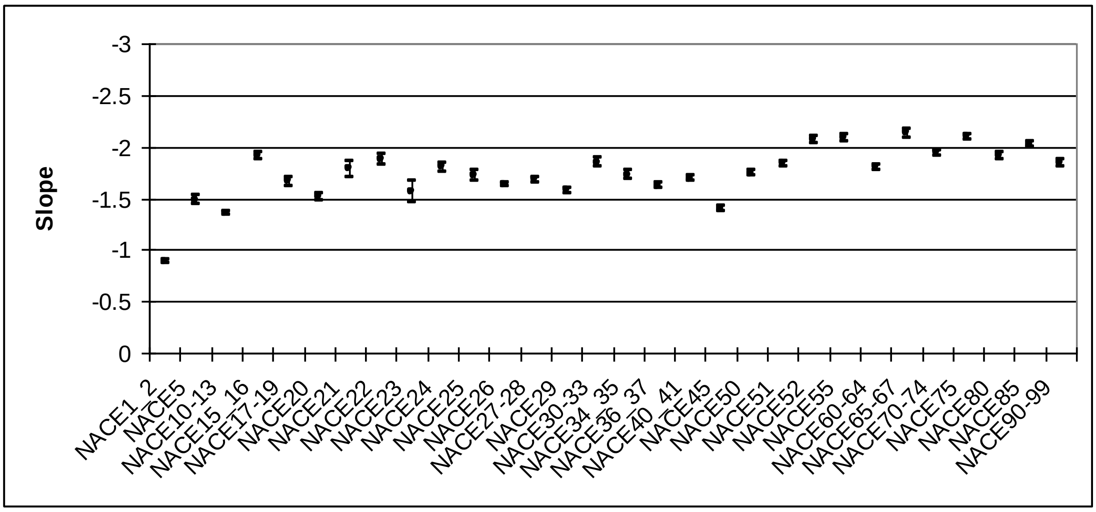

Source: Own calculations

Figure 5 Coefficient of the Urban Dummy

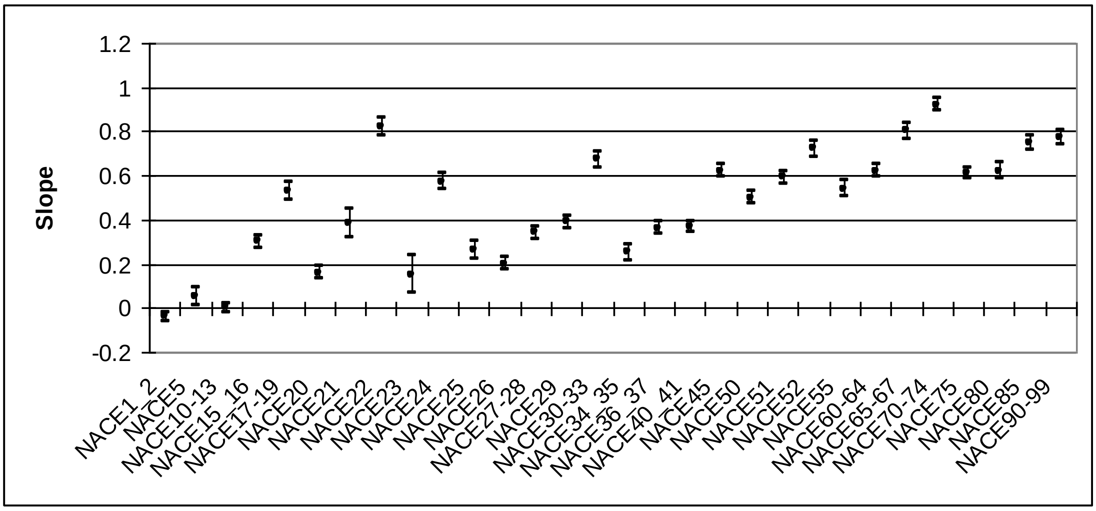

Source: Own calculations

6. SUMMARY AND CONCLUSIONS

This paper has addressed an important gap in the literature in that it has established the economic geography of Ireland for 2006 using a novel approach that was facilitated by the geocoding of the place of work by the CSO. In contrast to previous work this analysis was able to consider the economic geography at a sectorally highly disaggregated level, including a breakdown of the services sector, and at the spatially most disaggregated level.

The analysis confirms that the spatial distribution of employment differs significantly between sectors. The formal analysis confirms that this is not simply a chance outcome but that it is systematic and in many cases statistically significant. It highlights the spatial heterogeneity of the location of employment at the local level. Overall, the spatial heterogeneity is higher within the larger administrative units such as counties, regional authorities and regional assemblies that between these units.

The fact that there are statistically significant differences of employment location between sectors suggests that the locational requirements of the sectors differ. The paper considered just one underlying factor namely urbanisation. This analysis has shown the strong preference of certain sectors for urban locations.

Adam Smith (1776) in his seminal book already identified that some sectors will only be found in cities "There are some sorts of industries, even of the lowest kind, which can be carried on nowhere but in a great town. A porter, for example, can find employment and subsistence in no other place. A village is by far too narrow a sphere for him;…".

This has important implications for regional policy. If sectors have very specific locational requirements, and the analysis here suggests they do, then a policy of spreading employment will be counterproductive in a globalised world economy where firms are free to seek the most profitable location for their activities at a global level. Thus, while such a policy might reduce regional disparities, it is also likely to result in overall lower welfare.

Clearly, more research is necessary to uncover all the factors that drive the locational requirement of individual sectors in order to identify sensible policy measures. This analysis is left to future work, which is now possible given that the required data has been established in this paper.

As was identified in the introduction, the nature of the economic geography also has other implications. A high level of specialisation in low growth sectors will lead to low growth in these areas. On the other hand, a high level of specialisation in any sector, including high growth sectors, makes an individual area susceptible to shocks to the sector in which it is specialised. Thus, while financial services have grown substantially over the last decade or more, the recent financial crisis might impact negatively on those EDs that have a significant specialisation in that sector. Likewise, the construction sector is contracting rapidly at the moment, which will not impact equally across the country. The implications of such structural economic change could not be investigated as part of this paper but are analysed in a companion paper to this one (Morgenroth, 2008b).

References

Anselin, L., (1995) “Local Indicators of Spatial Association – LISA”, *Geographical Analysis,* 27: 93-115.

Bickenbach and E. Bode (2008) “Disproportionality Measures of Concentration, Specialization and Localization”, *International Regional Science Review,* 31: 359-388.

Breathnach , P. (2000) “The evolution of the spatial structure of the Irish dairy processing industry”. *Irish Geography,* 166-184

Cliff, A.D., and K. Ord, (1981), *Spatial Processes: Models and Applications*. London: Pion Limited.

Crampton, J W., (2001) “Maps as social constructions: power, communication and visualization”, *Progress in Human Geography*. 25:235-252.

Egeraat, C. van, and D. Jacobson. 2006. "Geography of Production Linkages in the Irish and Scottish Microcomputer Industry: the Role of Information Exchange". *Journal of Economic And Social Geography*, 97, 4, pp405-417.

Egeraat, C. van, and D. Jacobson. 2005. "Geography of Production Linkages in the Irish and Scottish Microcomputer Industry: the Role of Logistics". *Economic Geography*, 81, 3, pp283-303.

Gehlke, C. and Biehl, K. (1934) “Certain Effects of Grouping Upon the Size of the Correlation Coefficient in Census Tract Material”, *Journal of the American Statistical Association*, 29:169-170.

Geary, R C., (1954) “The contiguity ratio and statistical mapping”, *The Incorporated Statistician*, 5:115-145.

Gleeson, A-M., Ruane, F., and J Sutherland (2006) “Public Policy, Sectoral Specialisation and Spatial Concentration: Irish Manufacturing 1985-2002”, *Journal of the Statistical and Social Inquiry Society of Ireland*, 35: 110-150.

Krugman, P. A. (1991), “Increasing Returns and Economic Geography”, *Journal of Political Economy*, 99: 483-99.

Moran, P. A. P., (1948) “The interpretation of statistical maps” *Journal of the Royal Statistical Society B*, 10: 243-251.

Morgenroth, E., (2008a) “Economic Integration and Structural Change: The Case of Irish Regions” in Krieger-Boden, C., Morgenroth, E. and G. Petrakos (eds.) *The Impact of European Integration on Regional Structural Change and Cohesion.* London: Routledge.

Morgenroth, E., (2008b) “Simulating the Impact of Structural Economic Change on the Spatial Distribution of Job Location in Ireland” *ESRI Working Paper*

Paci, R., and F., Pigliaru (1999) “Is Dualism Still a Source of Convergence in Europe?”, *Applied Economics*, 31:1423-1436.

Smith, A., (1776) *The Wealth of Nations*

Strobl, E., (2004) “Trends and Determinants of Geographic Dispersion of Irish Manufacturing Activity, 1926 –1996”, *Regional Studies,* 38: 191-206.

Tobler, W., (1970) "A computer movie simulating urban growth in the Detroit region". *Economic Geography*, 46: 234-240.

Map 3 Job Density Agriculture and Forestry (persons per km^2), 2006

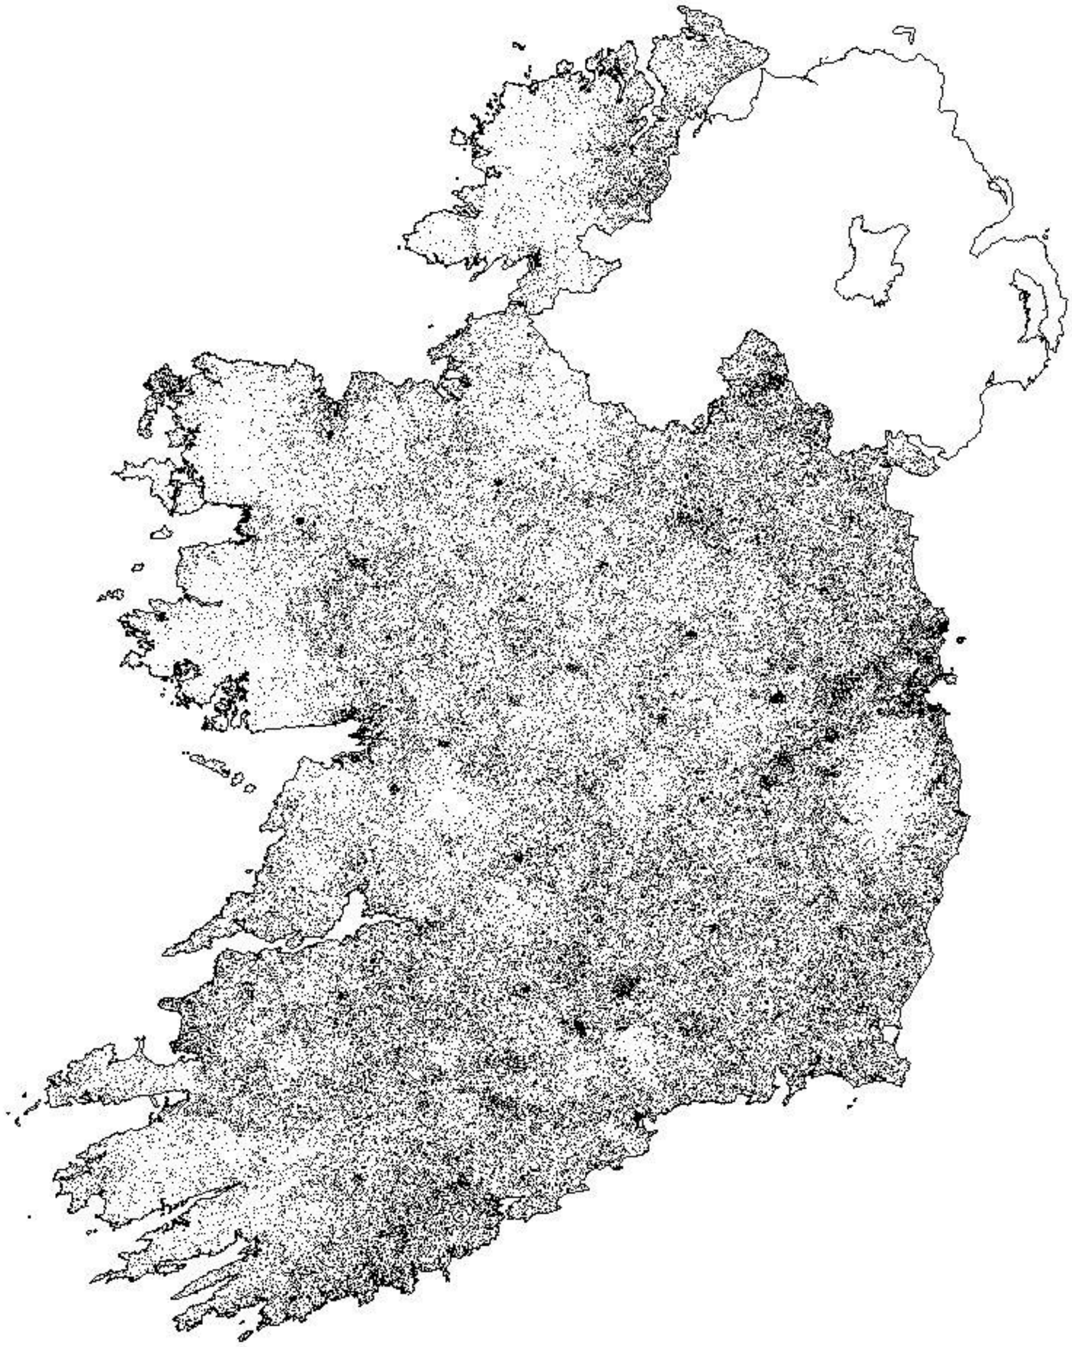

Map 4 Job Density Food and Drink (persons per km^2), 2006

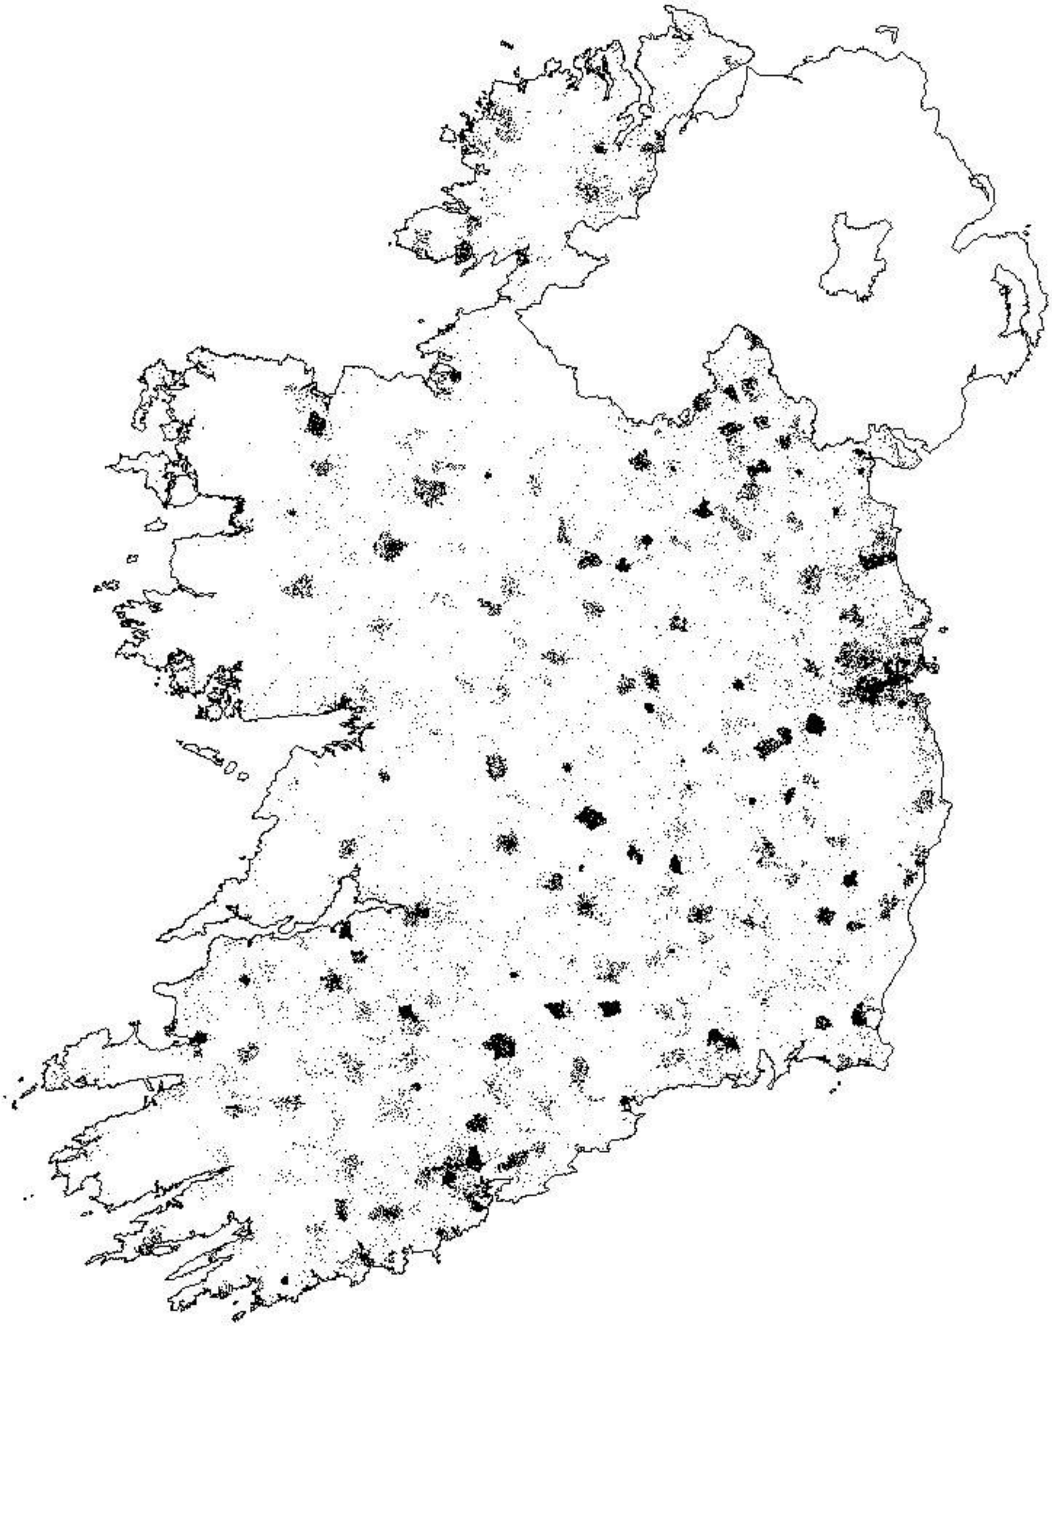

Map 5 Job Density Chemicals and Chemical Products (persons per km^2), 2006

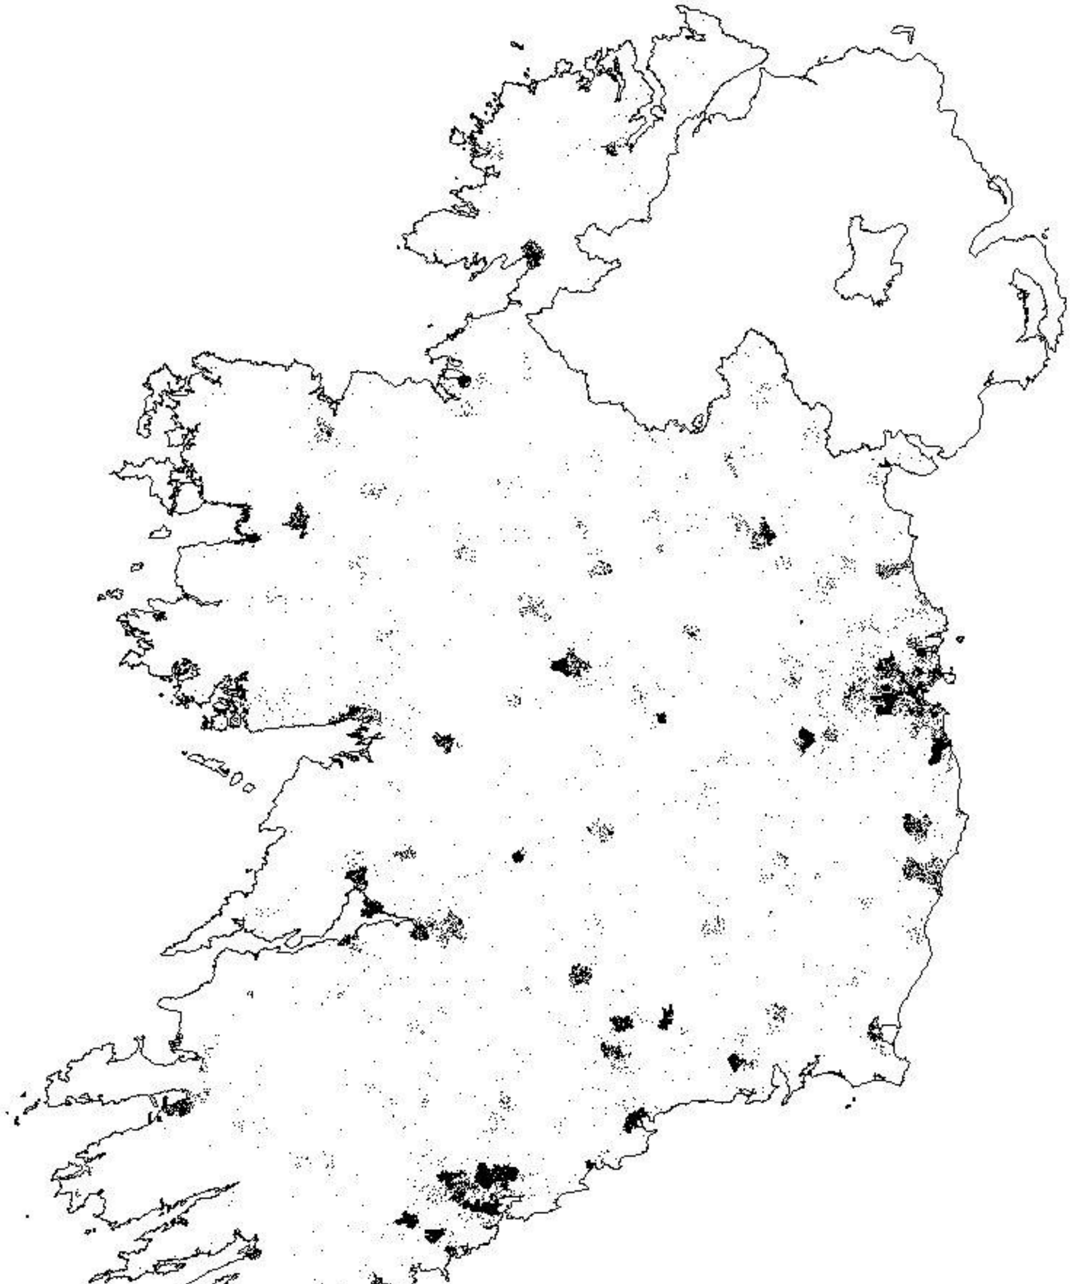

Map 6 Job Density Construction (persons per km^2), 2006

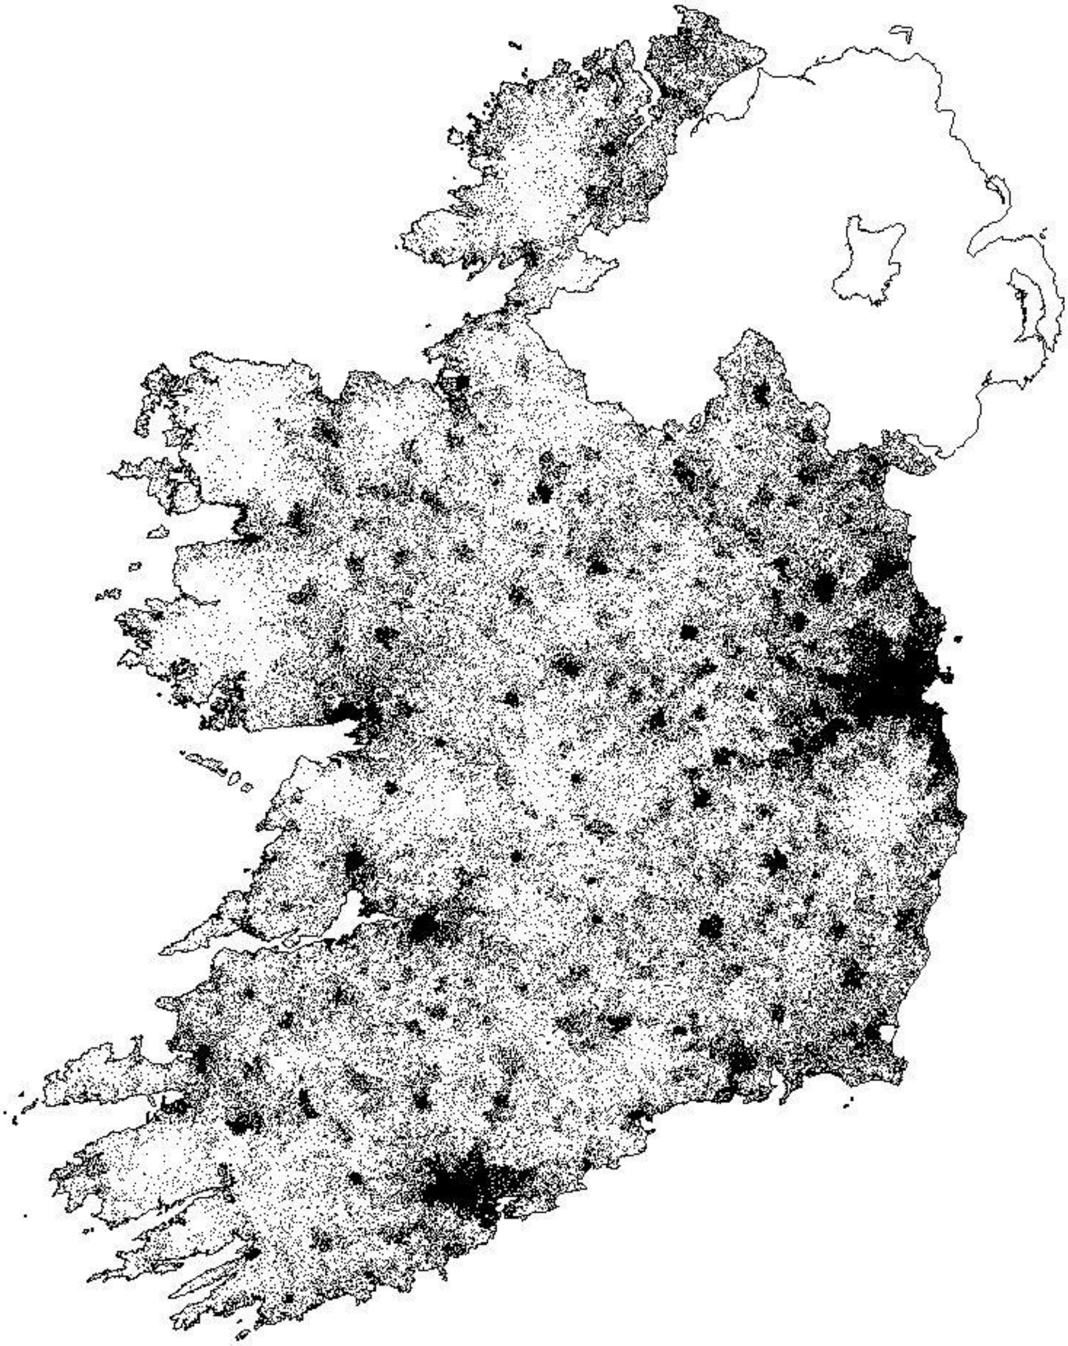

Map 7 Job Density Financial Services (persons per km^2), 2006

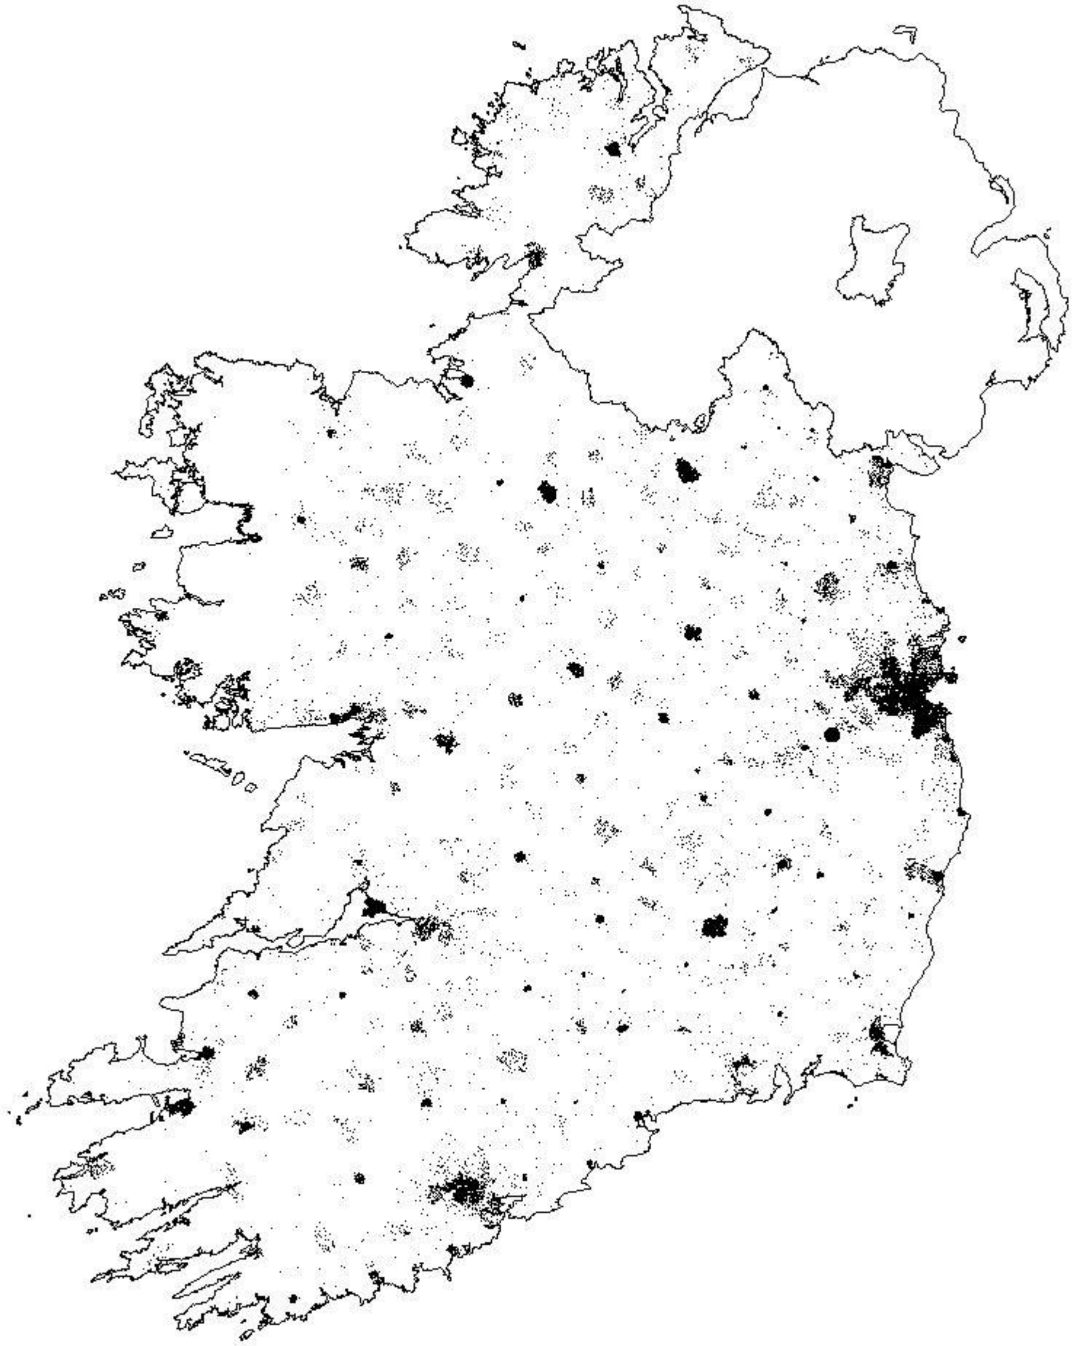

Map 8 Job Density Education (persons per km^2), 2006

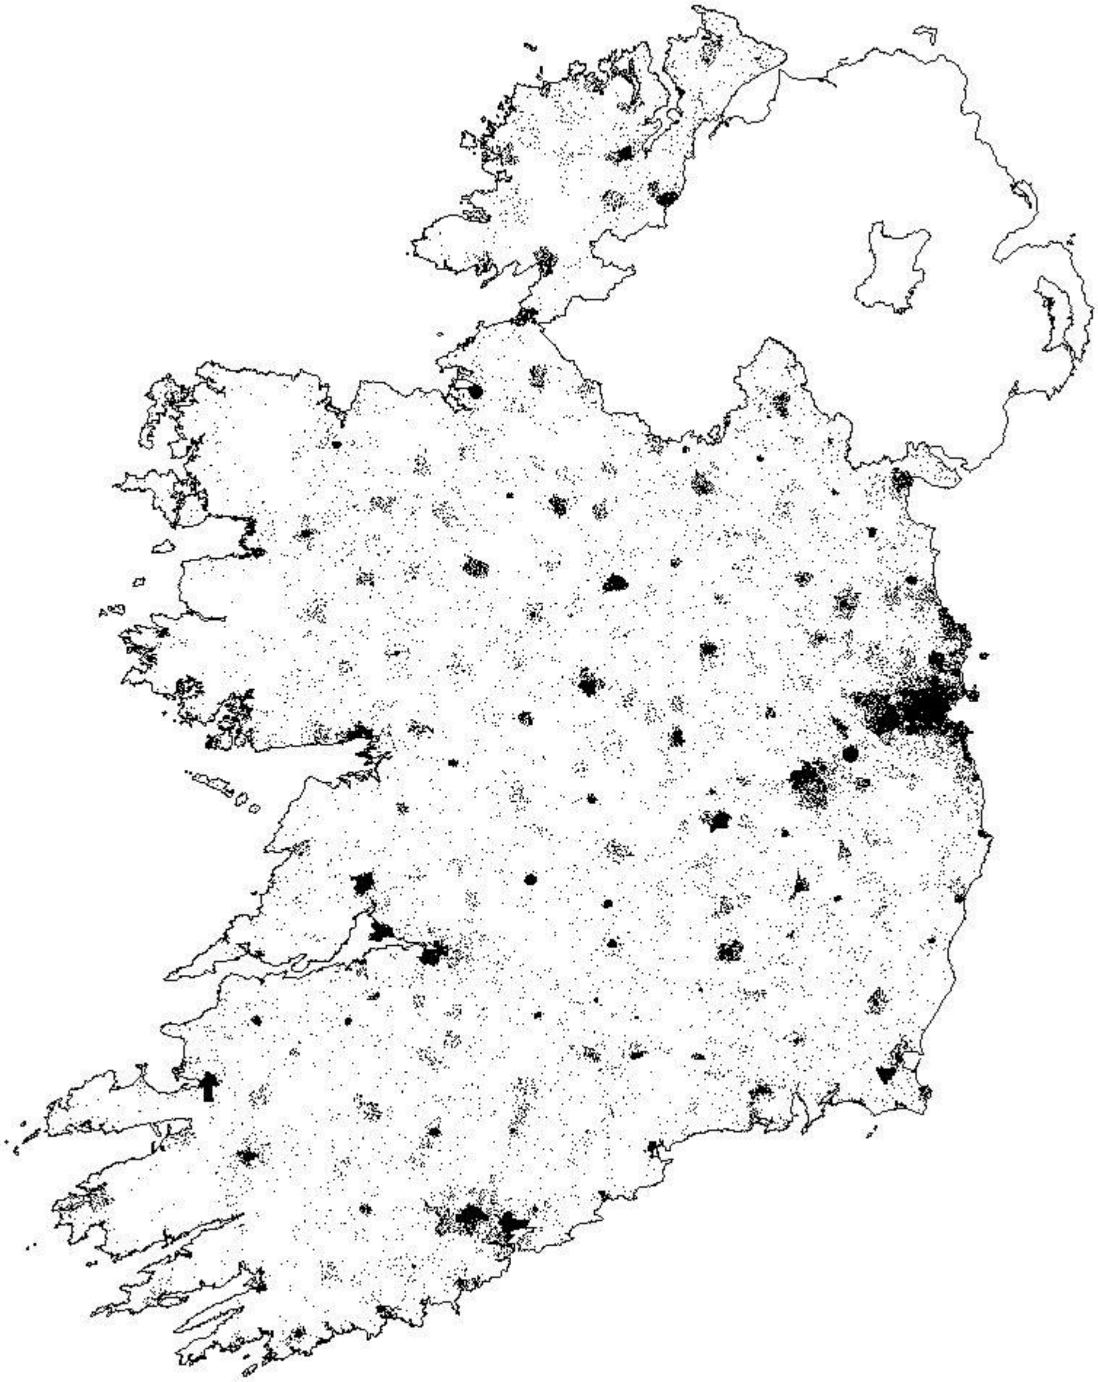

Map 10 Local Indicators of Spatial Autocorrelation (LISA) Map for Agriculture and Forestry

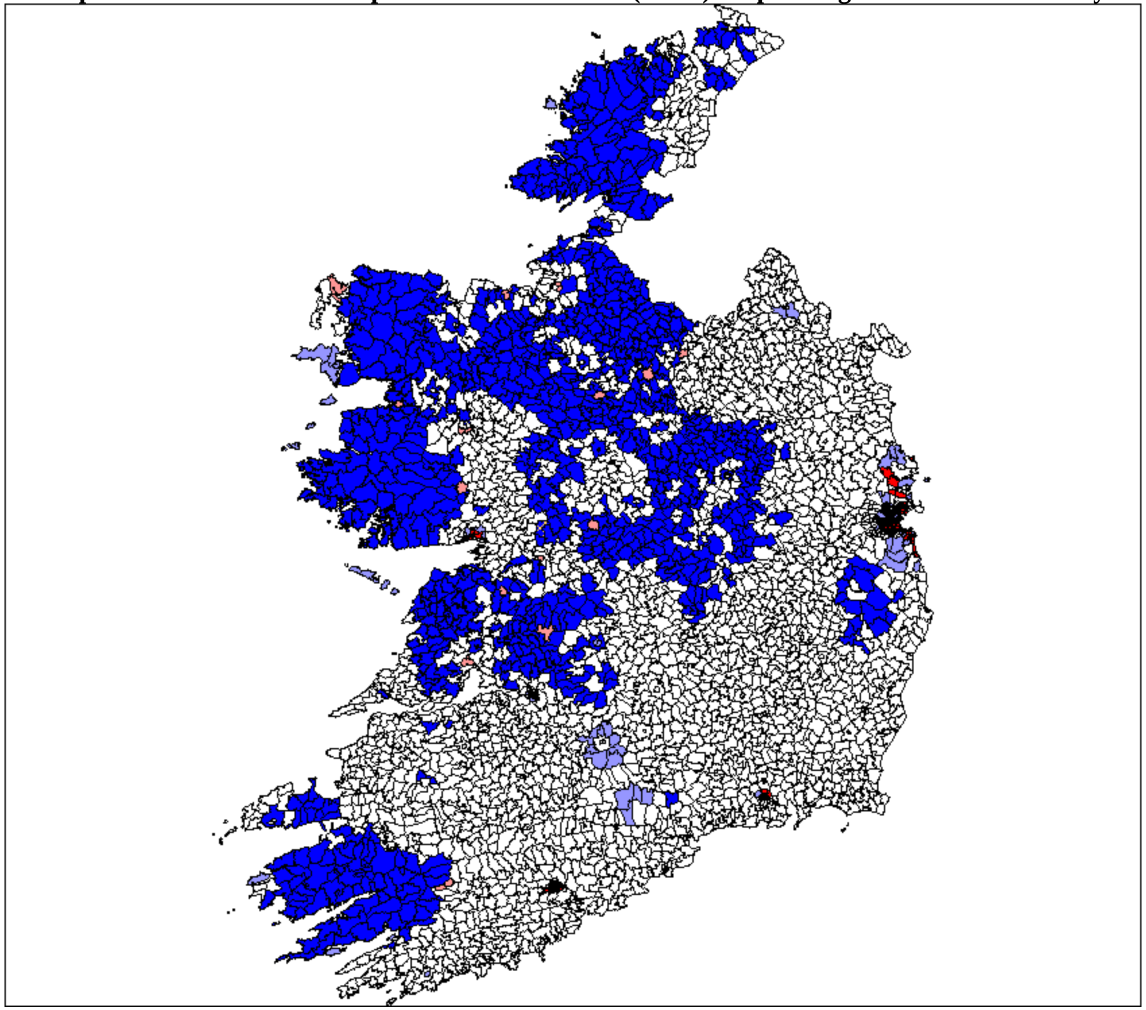

Note: White indicates LISA statistics that are not statistically significant, dark blue indicates local correlation between low density EDs, light blue indicates correlation between low and high densities, pink indicates a correlation between high and low densities and red indicates a correlation between high density EDs.

Map 11 Local Indicators of Spatial Autocorrelation (LISA) Map for Food and Drink

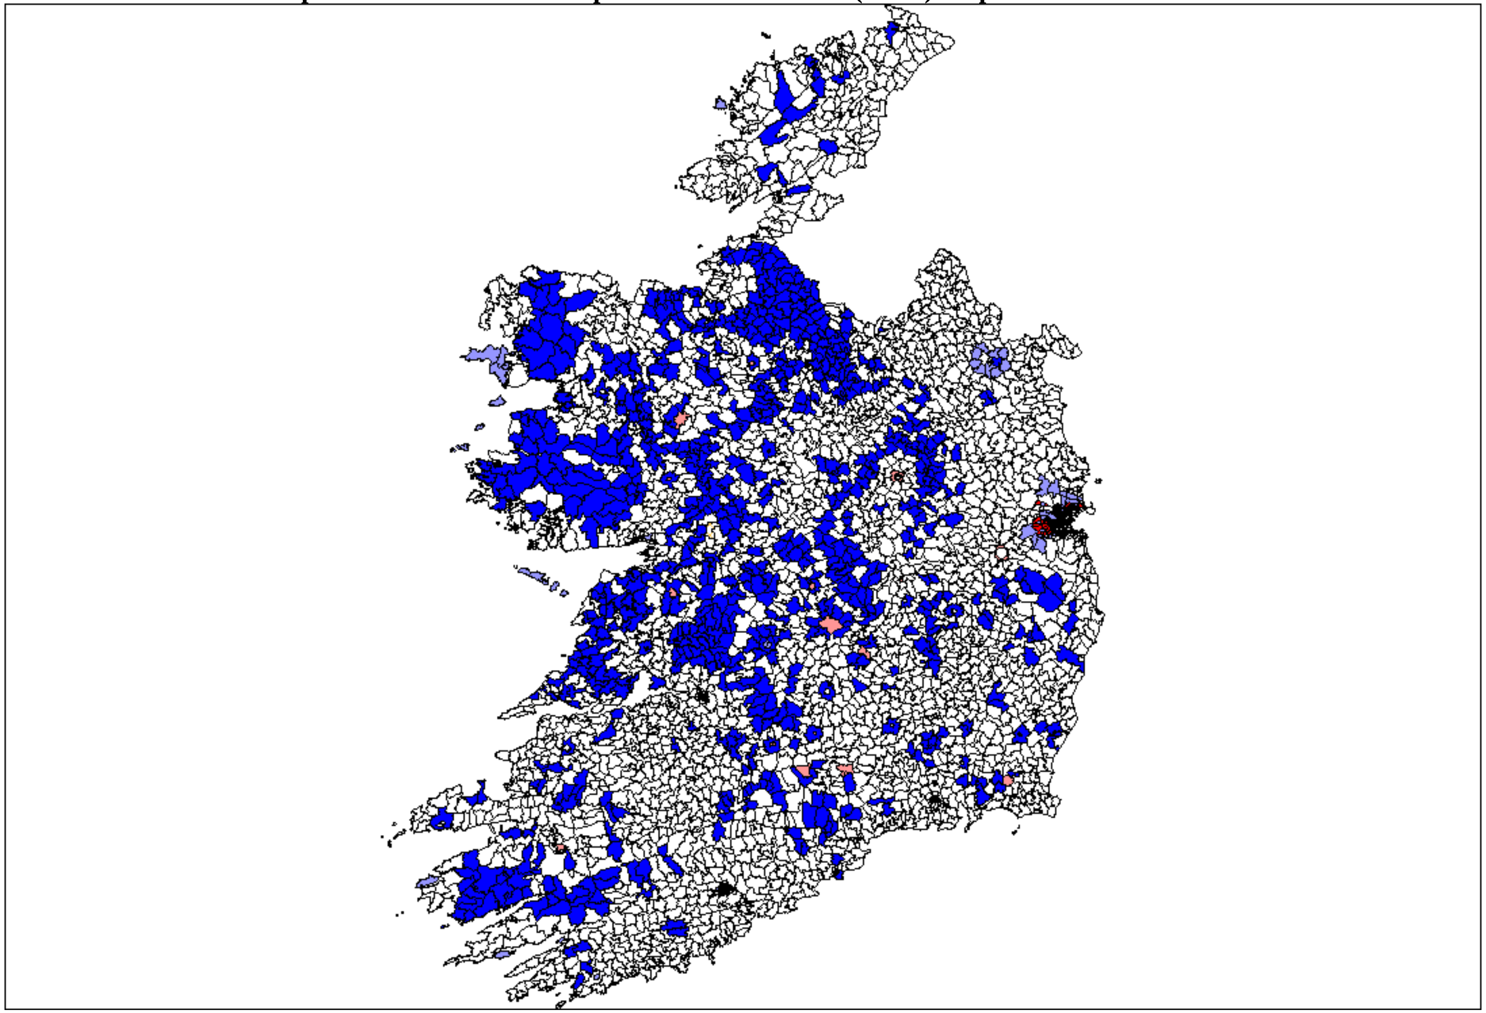

Note: White indicates LISA statistics that are not statistically significant, dark blue indicates local correlation between low density EDs, light blue indicates correlation between low and high densities, pink indicates a correlation between high and low densities and red indicates a correlation between high density EDs.

Map 12 Local Indicators of Spatial Autocorrelation (LISA) Map for Chemicals and Chemical Products

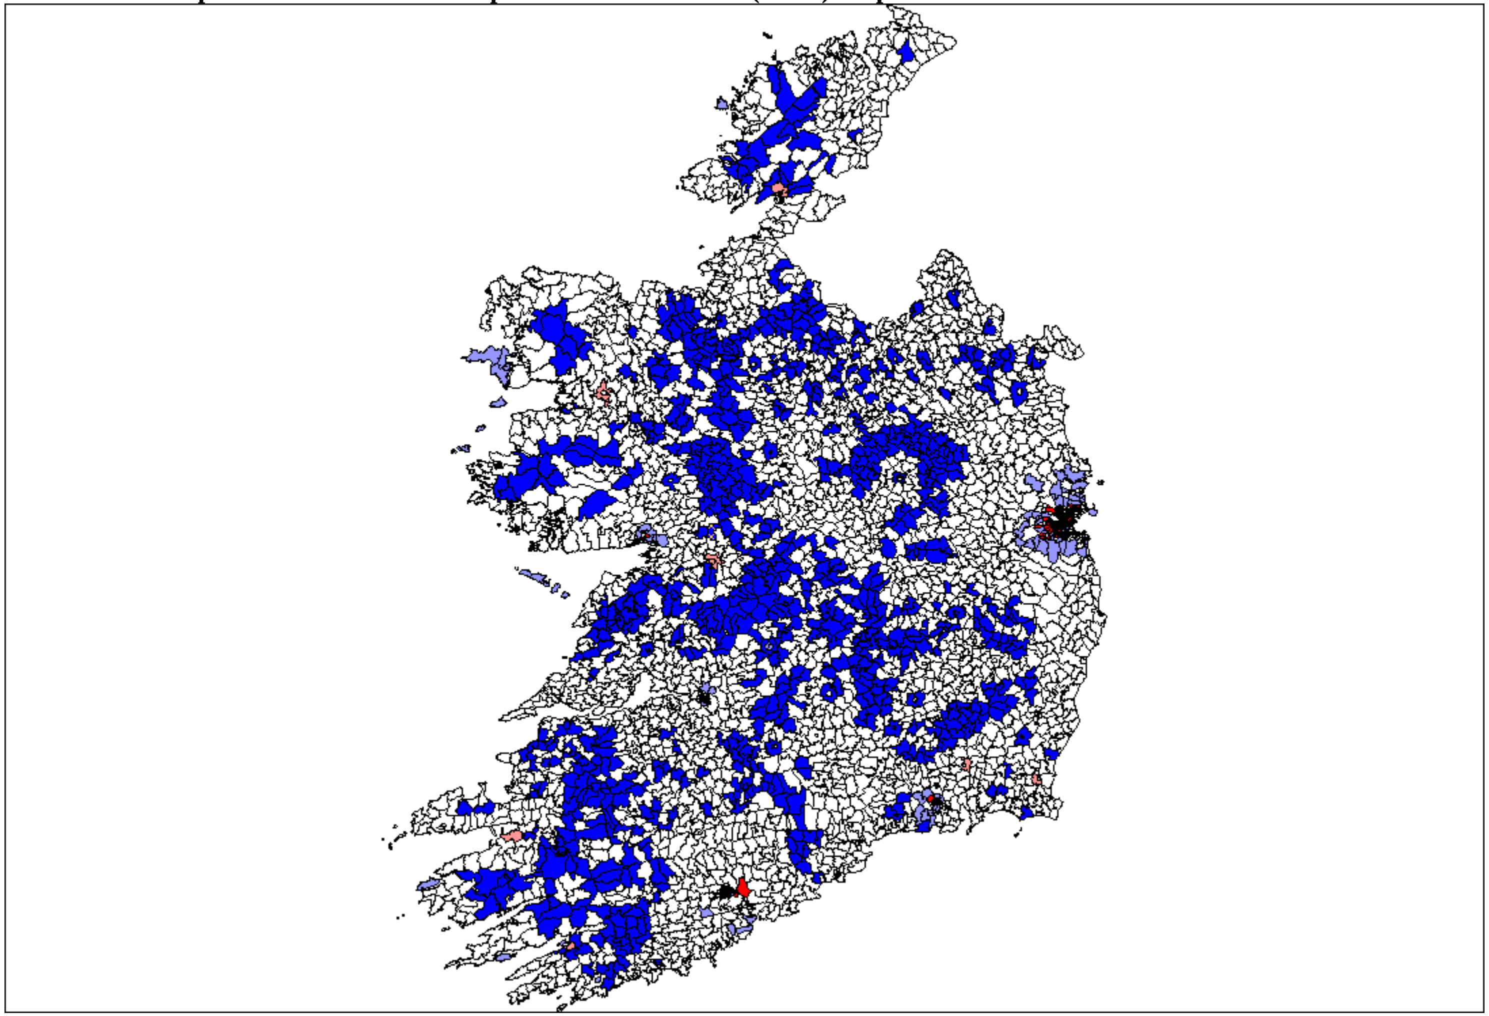

Note: White indicates LISA statistics that are not statistically significant, dark blue indicates local correlation between low density EDs, light blue indicates correlation between low and high densities, pink indicates a correlation between high and low densities and red indicates a correlation between high density EDs.

Map 13 Local Indicators of Spatial Autocorrelation (LISA) Map for Construction

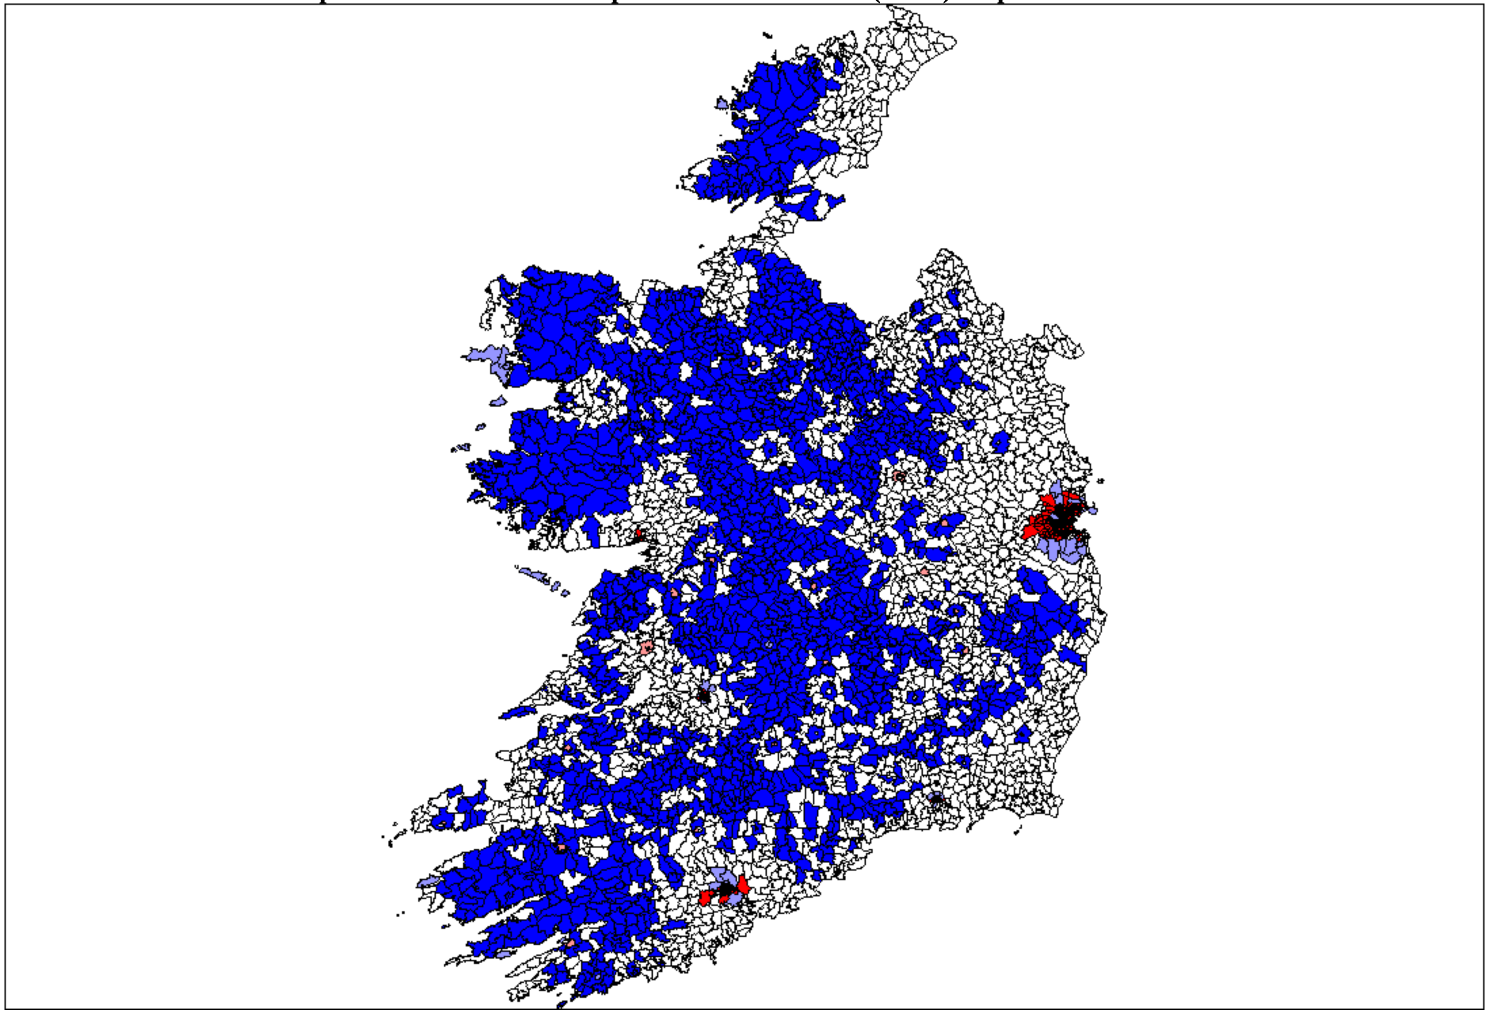

Note: White indicates LISA statistics that are not statistically significant, dark blue indicates local correlation between low density EDs, light blue indicates correlation between low and high densities, pink indicates a correlation between high and low densities and red indicates a correlation between high density EDs.

Map 14 Local Indicators of Spatial Autocorrelation (LISA) Map for Financial Services

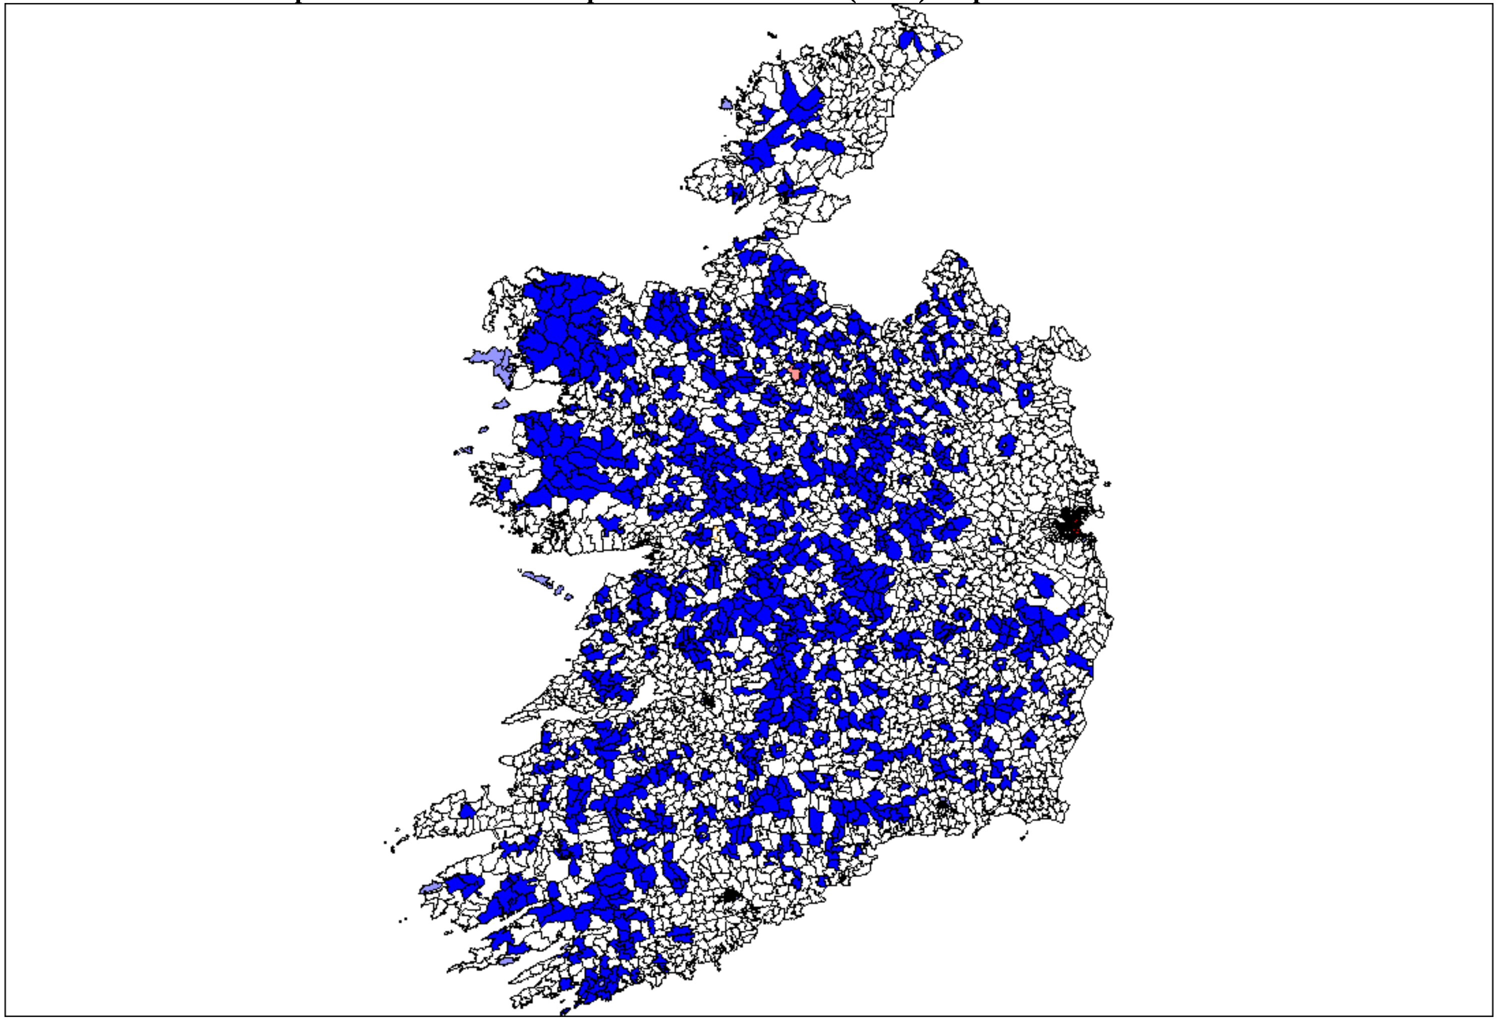

Note: White indicates LISA statistics that are not statistically significant, dark blue indicates local correlation between low density EDs, light blue indicates correlation between low and high densities, pink indicates a correlation between high and low densities and red indicates a correlation between high density EDs.

Map 15 Local Indicators of Spatial Autocorrelation (LISA) Map for Education

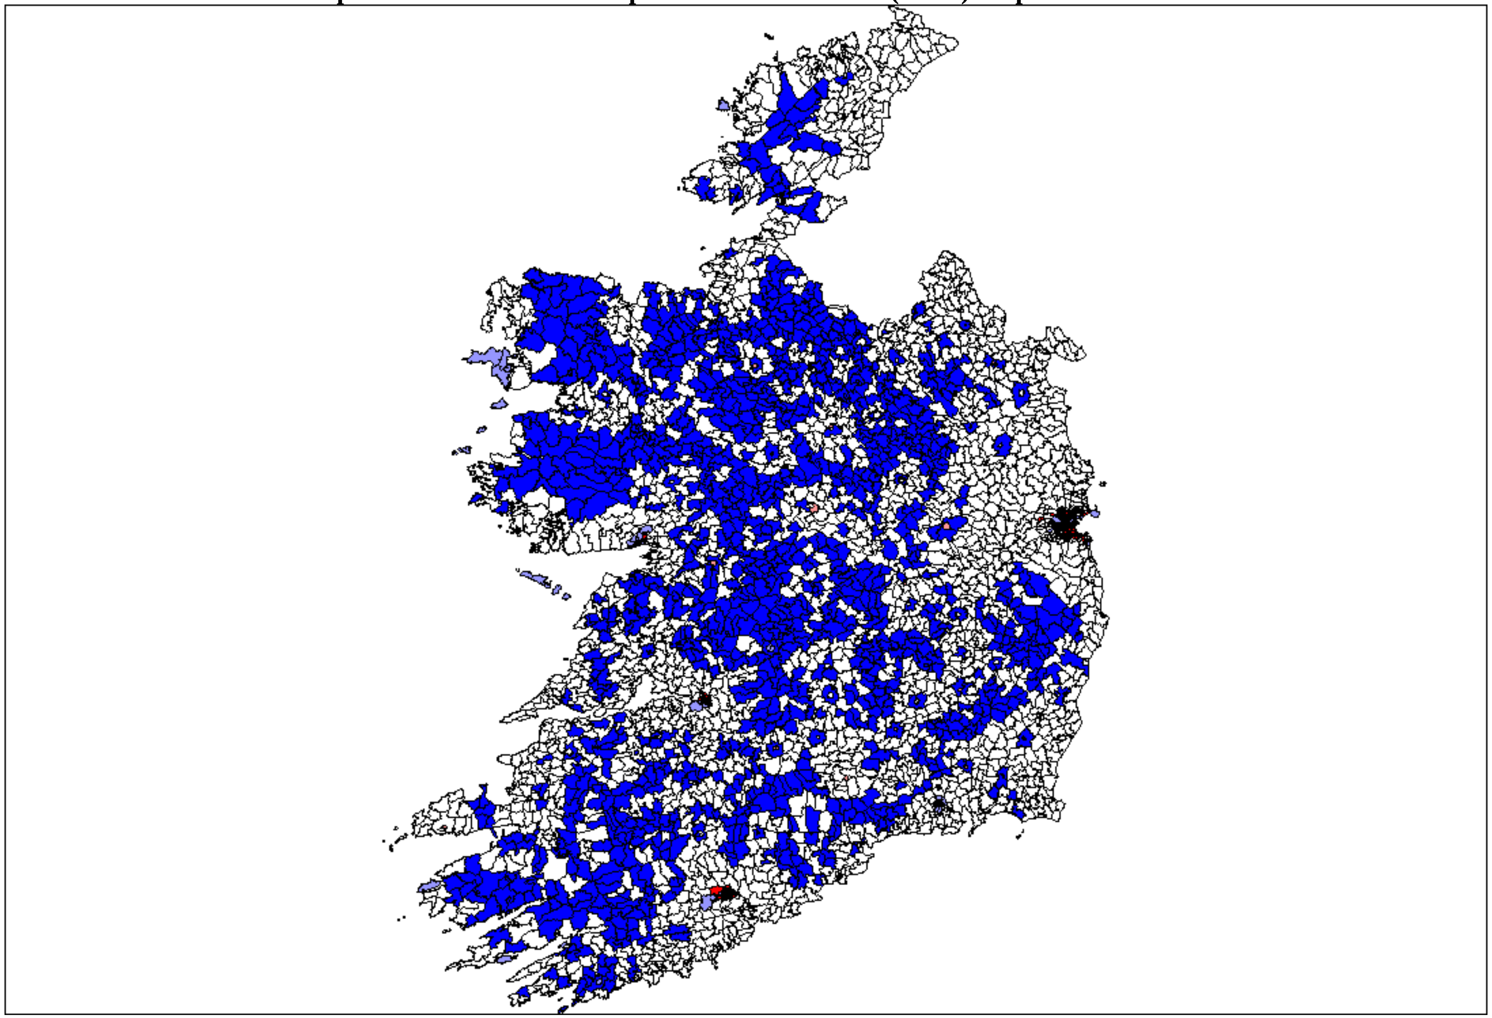

Note: White indicates LISA statistics that are not statistically significant, dark blue indicates local correlation between low density EDs, light blue indicates correlation between low and high densities, pink indicates a correlation between high and low densities and red indicates a correlation between high density EDs.

First vote of Thanks Proposed by Professor David Jacobson, Dublin City University

The admirable brevity of this paper contrasts remarkably with the effort that went into its production and even more with the contribution that it makes to the current knowledge and understanding that economics has of the economic geography of Ireland. The specification of new output from the CSO, and the clever use to which it has been put are just part of what Dr. Morgenroth has added. His deep understanding of regional economics and economic geography and of the statistical methods that can be used in their analysis have come together in this substantial contribution.

Given the title of the paper, a number of questions arise at the level of the discipline, or sub-discipline. First, one can ask what the differences are among "economic geography" (see such journals as *Economic Geography* and *Journal of Economic Geography*), "regional economics" (*Regional Studies* and *European Urban and Regional Studies*) and "regional science" (*Annals of Regional Science*, *Journal of Regional Science* and *International Regional Science Review*). At one time it may have been possible more clearly to distinguish among these disciplines. Economic geography consisted of the study of the distribution of economic activity, and in particular industrial activity, spatially within and between countries. Regional economics focused more on evenness or otherwise of activity – and income and welfare – between regions within economies. And regional science was distinguished more by the use of more sophisticated statistical methods in analysis of spatial aspects of economic activity than by the subject matter focused on by the experts in the area. Today, a perusal of the journals in these three areas will show that these distinctions are at the very least less sharp than they used to be. A paper like the one here produced by Dr. Morgenroth would appropriately be published in any of the major journals, including, among many others, those mentioned above.

This convergent tendency is a consequence of the increasing use – and increasing sophistication – of quantitative methods across these disciplines. This is not necessarily a bad thing if it adds to the accuracy and validity of the contributions that work in these disciplines can make, either to policy or to academic understanding. Few would argue that all the examples of increasing sophistication in the use of quantitative methods in these disciplines achieve these ends.

Leontief famously criticised theoretical model-building and, in his American Economics Association Presidential address in 1971, warned against the mathematical formalism already then dominating the profession. In essence, the target of his attacks was overly deductive theorising, and in particular the assumption of equilibrium underlying the models. Dr. Morgenroth's paper eschews such assumptions and focuses on the analysis of a new dataset. It is sophisticated and innovative in its use of statistical methods. For these reasons Leontief would probably have enjoyed Dr. Morgenroth's paper. It is highly inductive, analysing observations in a variety of ways, some of them new and unusual.

Despite this, the paper arrives at conclusions that could be argued to be obvious:

- "…the spatial distribution of employment differs significantly between sectors";
- "…locational requirements of … sectors differ";
- certain sectors have strong preference for urban locations;
- regional policy should not aim to spread employment evenly across the economy;
- an area highly specialised in low-growth sectors will itself have low growth; and,
- an area highly specialised in any one sector is susceptible to shocks to that sector.

These conclusions could well have been arrived at without the new dataset, and without the sophisticated quantitative methods utilised in this paper. This should not be taken as rejection of the value of the paper. It does, after all, lay groundwork in the utilisation of the new POWCAR data and more detailed and, perhaps, more policy-relevant conclusions can be expected to be forthcoming from subsequent analyses.

Nevertheless, there is a case to be made, even in the context of this, a statistical society, for space to be found in the journals for research papers using qualitative methodologies. This is all the more so, given the convergence in disciplines (or sub-disciplines), given the increasing predominance of quantitatively sophisticated papers, and given the nature of the conclusions of Dr. Morgenroth's paper. Approaches to the generation of knowledge other than those that have come to dominate, should not be ignored.

Leontief was correct in his 1971 warnings about the overly deductive turn in economics. His warnings were, however, not heeded. Leontief himself pointed this out in an interview in 1997 (Foley, 1998). In all economic research, including that with regional and geographic relevance, the trend has been towards more deductive theorising, and more quantitative methods. In its underlying – more inductive – theoretical methodology, Dr. Morgenroth's paper deviates from this trend and is to be commended for this.

Another deviation from the dominant trend is qualitative research. Let us consider arguments in favour of case studies. Case studies can be either quantitative and/or qualitative. Both are important. Quantitative methods are widely acceptable, qualitative increasingly less so. How can qualitative case studies contribute to knowledge in these areas? Most economists and regional scientists are sceptics when it comes to case studies and it is therefore worth presenting in some detail the answers of some of those who have used and justified qualitative case studies.

In my own work and that of my PhD students and associates over the past 15 years we have undertaken studies of a range of industries that would concur with many of Dr. Morgenroth's conclusions. For example, from my study of the car industry (Jacobson, 1977) it is clear that over-dependence on an industry is risky. When Ford ceased production – later moving to assembly – at Cork in 1932, employment was reduced from a high of 7,000 to 300 (*Irish Times*, 8/11/1932, p.7). In the midst of depression, this had a huge impact on the city. The current economic crisis provides fresher examples. Subsequent work on industrial location (Andreosso and Jacobson, 2005, Ch.8) has shown clearly that different returns to scale in different industries result in different levels of concentration. The unevenness of industrial activity is documented in my work with McDonough (Jacobson and McDonough, 1999) in what we called the 'lumpiness' of spatial concentration of industrial activity. Similarly other cases – and comparisons between them – provide evidence of others of the conclusions.

But can one generalise from case studies? Flyvbjerg (2006) addresses this and related misconceptions of qualitative research. He shows how and where generalising from even single cases is relevant and appropriate. His examples include Galileo's rejection of Aristotle's theory of gravity, on the basis of a single experiment, and the generalisation from cases in the physics of Newton and Einstein and in the evolutionary theory of Darwin. In relation to his own work, Flyvbjerg (2006) explains that having been trained as a neoclassical economist to expect effective markets, he undertook an in-depth case study of Aalborg, only to discover that:

> Members of the local business community were power mongers who were busy negotiating illicit deals with politicians and administrators on how to block competition and the free market and create special privileges for themselves. The neoclassical model was effectively falsified by what I saw in Aalborg.

His conclusion on the question of generalisability from cases is as follows:

> One can often generalize on the basis of a single case, and the case study maybe central to scientific development via generalization as supplement or alternative to other methods. But formal generalization is overvalued as a source of scientific development, whereas "the force of example" is underestimated.

Dr. Morgenroth writes that "more research is necessary to uncover all the factors that drive the locational requirement of individual sectors in order to identify sensible policy measures". At very high levels of industrial disaggregation it is likely that the drivers of location of a particular sub-sector of less than ten firms are different from those of a different sub-sector of similar size, even from the same industry. The more disaggregated the different sub-sectors, and the fewer the number of firms included in these sub-sectors, the more appropriate the case study method of research. And the more disaggregated the information, the more will in-depth research of those sub-sectors be necessary for policy advice.

This is clearly not to say that qualitative research is better or more important than quantitative research, only that it should not be ignored. Dr. Morgenroth himself would agree with this. His paper may indeed lead to the identification of appropriate case studies. It more than deserves the vote of thanks.

References:

Andreosso, B. and D. Jacobson (2005) *Industrial Economics and Organization: A European Perspective*, McGraw-Hill, UK.

Flyvbjerg, B. (2006) "Five Misunderstandings about Case-Study Research", *Qualitative Inquiry*, Vol. 12, No. 1, pp.219-245.

Foley, D. K. (1998) "An Interview with Wassily Leontief", *Macroeconomic Dynamics*, Vol. 2, pp.116-140.

Jacobson, D. S. (1977) "The Political Economy of Industrial Location: The Ford Motor Company at Cork 1912-26", *Irish Economic and Social History*, Vol.4, pp.36-55.

Jacobson, D. and T. McDonough (1999) "International Trade and European Integration", in E. O'Shea and M. Keane (eds.), *Core Issues in European Integration*, Oak Tree Press, Dublin.

SECOND VOTE OF THANKS PROPOSED BY PATRICIA O'HARA, POLICY MANAGER, WESTERN DEVELOPMENT COMMISSION

I am pleased to second this vote of thanks to Dr Morgenroth who has made such a significant contribution to the revitalisation of regional studies in Ireland in recent years. I would like to congratulate him on producing a pioneering and interesting paper which provides new insight into some critical regional development issues in Ireland.

As is pointed out in his introduction, there has been relatively little attention paid to the spatial distribution of economic activity in Ireland mainly due to the lack of data at detailed geographical levels. For this reason, the decision by the Central Statistics Office to release a dataset known as the Place of Work Census of Anonymised Records (PoWCAR) from the 2006 Census of Population is to be greatly welcomed. This dataset, by providing a complete count of the working population, geo-coded by place of work and residence at electoral division (ED) level, opens up the possibility of analysing, for example, the spatial distribution of employment, commuting patterns and labour catchments to throw light on the dynamics of local labour markets. Developments in geo-computation have facilitated the production of maps based on these data that can greatly enhance our understanding of the geography of local economies and assist in the planning of services.

The Western Development Commission is currently involved in a study of travel to work patterns and identification of local labour catchments which draws on the PoWCAR dataset. In this exercise we have also encountered some of the technical issues that Dr Morgenroth refers to in the paper, such as how to treat mobile workers, or those who did not report a workplace, both of which are excluded from the dataset. We have noted that that construction and transport workers are strongly represented within the mobile worker category, and that farmers are over-represented among those for whom no workplace was coded. I note that in this paper mobile workers are attached to their home residence and while this may be adequate for some purposes including this paper, it is more problematic when trying to identify labour catchments and travel to work patterns.

Turning to the substance of the paper, the patterns in Maps 1-2 are not surprising. Jobs are more clustered than population, and are concentrated in urban centres. Maps 3-8 show that jobs based on natural resources such as agriculture and forestry are very widely dispersed. However, there is also evidence of the concentration of farming employment in the stronger farming areas of the east and south, a tendency that has become more pronounced in recent decades. This is also visible in the map of employment in the Food and Drink sector where the concentrations are such as to allow us to associate job densities with the impact of particular plants.

The extent of the dispersal of construction employment in 2006 is very evident from Map 6, although this is a pattern which must have altered considerably since the economic downturn. From a regional perspective, what is most striking is, not the concentration in the largest centres which is to be expected, but the extent of the spread of construction jobs across all regions. In the West, some of this is associated with the Upper Shannon Tax Incentive Scheme which stimulated a construction boom in its catchment during the last decade. In the seven western counties that make up the Western Region[*], one in six of all workers in 2007 were employed in construction, accounting for 15% of all employment. This share has fallen to 13% in 2008 and will undoubtedly continue.

The limited spread of the Financial and Other Businesses sectors beyond Dublin in Map 7 is also striking. Jobs in this sector, which are generally higher skilled and higher paid, are noticeably scarce in the regions. While accounting for the largest share of employment in the state generally, the sector is highly spatially concentrated. It is the ranked only sixth in terms of employment share in the Western Region, for instance. The spatial pattern

[*] The Western Region is the operational area of the Western Development Commission and comprises counties Donegal, Sligo, Leitrim, Roscommon, Mayo, Galway and Clare.

of employment in this sector is, of course, strongly associated with the deliberate decision to locate the International Financial Services Centre (IFSC) in Dublin.

The spatial pattern of jobs in education (Map 8) is also notable as it illustrates the dispersed pattern of public sector jobs and underlines the strong dependence on this sector in the regions. (I suspect that the pattern of jobs in the health sector would be similar). We are very conscious of this in the Western Region. In 2008, almost one quarter (24%) of jobs were in the public service compared to 22% in the rest of the state. Women are particularly dependent on this employment: 42% of them worked in the public service in 2008, compared to 36% in the rest of the state.

The results of the formal testing of the spatial distributions of employment are pretty much as expected i.e. larger sectors generally have a wider spread; it is not surprising that smaller sectors are more spatially confined. Indeed, the location of the more specialised sectors is often the result of deliberate policy decisions regarding location, as in the cases of financial services, electricity gas and water supply, airports and other transport facilities.

There are some statements which would have benefited from clarification and elaboration. For instance: 'EDs surrounding the major urban centres are the least specialised while some urban EDs are very specialised and many rural EDs have either a high or medium level of specialisation' *(p.11);* or, on p.13, the statement: 'The most autocorrelated sector is Construction, while the least is Manufacture of Transport Equipment. This suggests an interesting relationship between the measure of spatial concentration and that of spatial autocorrelation, in that the least concentrated sectors appear to be the most spatially correlated.' How are we to interpret these patterns? What do they suggest about the underlying relationship between sectors and their spatial patterns?

I find it difficult to be comment on Maps 10-15 without seeing them in more detail. However, I would question whether we can really attribute the patterns in Dublin to 'urbanisation economies' since one of the defining characteristic of a city is itself the existence of concentration. It is not easy to distinguish cause and effect here except where, as I have already mentioned, location is the result of deliberate policy choice.

Finally, I would like to again congratulate Dr Morgenroth for opening up a huge area of debate in this interesting and provocative paper. It is clear from the analysis that the spatial distribution of employment across economic sectors is related to scale, to locational decisions of government, and to the existence of population clusters, particularly if the sector in question is dependent on local demand and/or specialised. It is true, of course, that some sectors are inherently urban and require population concentrations to support them – even if the presumed relationship between buying power and location is being undermined by modern communications technology. It is less easy to see the links between the spatial distribution of sectors and the locational requirements of sectors. This is an area crying out for empirical analysis, particularly the impact of infrastructure availability on location decisions. There is also need for a thorough assessment of location policies and their impact. As regards further analysis of the PoWCAR dataset, it would be interesting to explore the links between the sectoral distributions and local labour markets and their gender, age and education profiles.