The Geographical Distribution of Poverty in Ireland

D.G. Pringle

Dept. of Geography and National Institute for Regional and Spatial Analysis, National University of Ireland, Maynooth.

Abstract

Government investment is often spatially targeted or tends to have a differential spatial impact. We require objective indicators to identify which areas have greatest need of assistance. We also need to be able to monitor the impact of various strategies on poverty over time. There are numerous problems associated with measuring poverty. There are also numerous problems associated with identifying its spatial dimension. These problems could be overcome, thereby permitting a more rational area-based approach to tackling the structural problems which cause deprivation to be more concentrated in certain areas. It is argued that area-based schemes should be diverse in nature and take advantage of local initiative and expertise.

Introduction

When invited to make a presentation on the geographical distribution of poverty in Ireland, my immediate instinct was to produce a map showing the distribution of poverty which could form the focus for a discussion regarding causal factors and possible policy implications. However, a little reflection indicated that there was very little that I could definitively say about the distribution of poverty in Ireland. In my defence, I would add that I believe that there is very little that anyone else can definitively say. The problem is not just a question of insufficient information, but rather relates to the elusive nature of what we mean by 'poverty' and the ways in which it is measured.

I will begin by outlining some reasons why I think it is important to have an objective analysis of the geographical distribution of poverty. I will then outline some of the problems which militate against producing such an analysis. Some of these problems are associated with the way in which poverty is measured, whereas other problems arise from the choice of spatial units for which these measurements are made. Both types of problem are reviewed. The paper concludes with a few observations about the type of strategy required to tackle the structural problems found in areas of deprivation.

The Need for an Objective Analysis

Government investment in development is often spatially targeted (i.e. it is intended to benefit specific identified areas). Even if the investment is not specifically targeted, it tends to have a differential spatial impact (i.e. it benefits some areas more than others just by its nature). When drawing up plans for investment, we should have some notion of its likely spatial implications. If the objective is to reduce poverty and to promote social inclusion, then clearly we need to have some idea of which areas have the greatest need so that we can target these areas for special attention.

We each have our own mental image of which parts of the country are the richest and the poorest. The first point I would make is that our perceptions can in some cases be very misleading. To illustrate, consider this map of the mean annual infant mortality rate for Irish counties for the period from 1916-1935 (Figure 1).[16] Infant deaths (i.e. deaths below the age of 12 months) generally reflect poor living conditions, affecting either the mother during pregnancy or the child after birth. Internationally present day infant mortality rates are highly correlated with measures of economic development, to the extent that they probably provide a more reliable indication of living standards than many economic indicators (which are often subject to so many cross-cultural qualifications as to be virtually useless). Infant mortality in Ireland in the past may therefore be regarded as a plausible indicator of living standards.
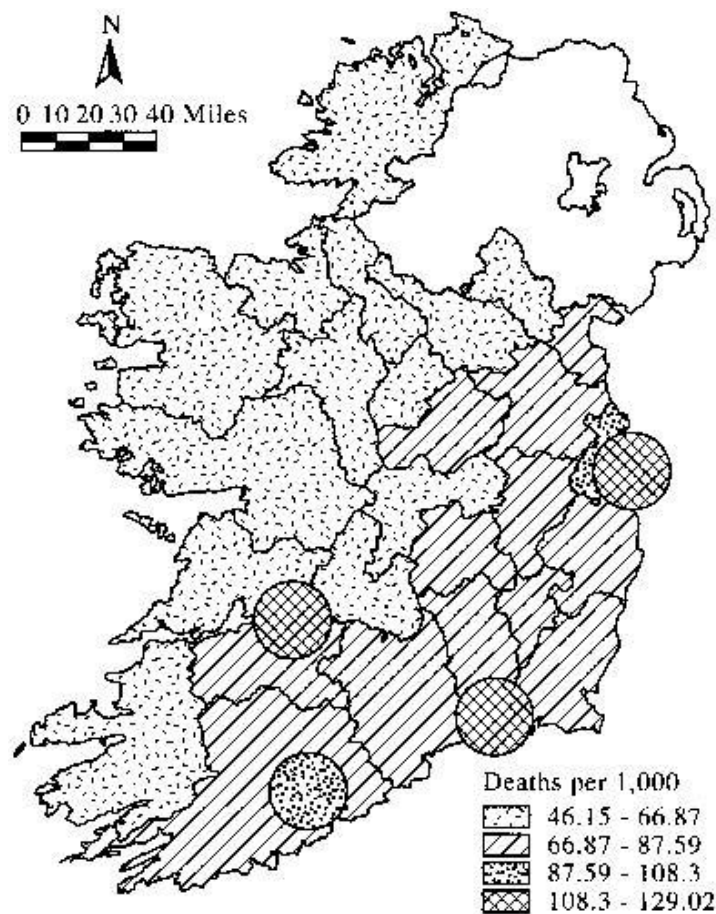
Figure 1. Mean Annual Male Infant Mortality Rates, 1916-35.

The first thing that one might note from the map is that there is a very distinct regional divide between the north and west of the country, and the south and east – corresponding with the major regional divide that pervades so many aspects of Irish life. The major cities (represented by the circles) form a separate category. The second point that might be noted is that there is a massive disparity between the counties with the lowest rates and those with the highest rates. We are not talking

---

[16] This map was originally drawn for a totally unrelated purpose. See Pringle (1998) for details.

about a few percentage points in the difference: the worst counties had a rate three times higher than the best counties. It should also be noted that the map is not a statistical aberration caused by small numbers. It is based on almost 50,000 infant deaths over a 20 year period. The same pattern was repeated, with minor fluctuations, year after year.[17]

The main point I would make about this map is that the pattern of infant mortality is quite different to what most people would have expected. The major urban areas had a higher rate of infant mortality than the counties. This runs counter to the notion of poverty in Ireland being a predominantly rural phenomenon, but might not be too surprising if one bears in mind the contemporary accounts of slum conditions in the major cities. However, what I think would surprise most people is that the infant mortality rates were higher in the more affluent south and east of the country than the north and west. In other words, infants had the best survival chances in those areas where most people would have expected them to be the worst.

The pattern of infant mortality in Ireland in the earlier decades of the last century clearly raises a lot of questions. I do not want to get into a discussion of the details here. The point I want to make is that if the empirical data had not been collected and mapped it most of us would have jumped to completely the wrong conclusions based upon our prior expectations. Likewise, rather than assume that we know which parts of the country are the most disadvantaged today, what we need to do is to collect hard information on poverty and map it.

It might be noted parenthetically that the geographical pattern of infant mortality today is quite different in two important respects. The first is that the overall rate is happily only a tiny fraction of what it once was and is influenced by factors other than living conditions to a much greater extent. The second is that because the number of infant deaths is small, the incidence rates no longer exhibit a clear geographical pattern and also tend to vary from one year to the next. Infant mortality can no longer be regarded as a reliable indicator of living conditions, even at a county level.

An objective analysis of the geographical distribution of poverty would enable us to identify which areas should be targeted for special attention. However, we also require an objective analysis to enable us to monitor the effectiveness of our policies. This aspect is all too often overlooked. Funds are often channelled into projects to benefit areas of identified need without any objective assessment of the impact of the project – i.e. we do not have any objective indication of whether things have got better or worse. More importantly, even if things have got better, we are often unable to say whether they got better <u>because</u> of the project or whether they would have improved anyway. In other words, we need an objective method for assessing the effectiveness of a given policy or project. Comparing the extent to which things have improved in an area selected for special attention compared with similar or adjoining areas not selected for special attention would provide one indication of the effectiveness of the project or policy. Also, if the objective is to promote social inclusion, then our policies should result in a reduction in spatial disparities. We need

---

[17] This analysis is based solely on male infant deaths. The patterns for females were almost identical, although the rates were slightly lower (approximately 38,000 deaths).

an objective method of assessing whether disparities are in fact increasing or declining.

To enable this type of monitoring we require objective measures which a) facilitate comparisons between areas at a given point in time (to enable us to identify which areas are better or worse, and by how much), and b) facilitate comparisons within each area for different points in time (so that we can establish whether things are getting better or worse, and by how much).

Measuring Poverty

Quantifying the amount of poverty in a given area is more complicated than one might assume. Most people would regard poverty as the opposite of wealth – i.e. it is a financial concept indicating a shortfall in disposable income relative to what is required for an acceptable standard of living. It should therefore ideally be measured by some sort of financial indicator. The problem is that there is no readily available source of data to do this. The Census of Population, which is the major source of information on many aspects of life, does not contain any questions on income. The Living in Ireland survey, conducted by the Economic and Social Research Institute, does contain income information, but although it is based on a very large sample by social survey standards (n=7,000), this translates to an average of only 2 households per DED (the smallest area for which statistics are routinely collected), which is clearly insufficient for a geographical analysis.[18] Income information is collected by other state agencies (e.g. Revenue Commissioners) for taxation purposes. This could conceivably be used to calculate an income indicator for small areas, but numerous logistical issues would need to be resolved before this would be feasible.

In the absence of an income measure, poverty indicators are normally based on indirect measures. In the absence of single ideal surrogate measure, the usual procedure is to use a number of surrogate indicators (e.g. unemployment, poor education, etc.) which are then combined to form a single composite index. A positive spin is often placed upon these procedures by noting that poverty (and related concepts such as multiple need, multiple deprivation, social exclusion, etc.) are 'multidimensional' and that this approach allows this multidimensional nature to be captured. In other words, multivariate measures are made to sound as if they are a good thing, rather than being what they really are - a substitute for good information (i.e. income measures).

The use of multivariate indices introduces a substantial degree of subjectivity into the analysis – i.e. different composite indices will give different results depending upon which surrogate measures of poverty are included and also upon the method used to combine them into a single composite measure. A number of composite indices have been calculated for Ireland, but the best known are probably the Haase index (Haase, 1995; 1999) and the SAHRU index (SAHRU, 1997). These each identify quite different areas as being the most deprived.

There is no agreement as to which indicators should be included in a composite index. Given that the Census of Population provides the major source of information, most indicators are drawn from the information available from the Census. The

---

[18] For an informative analysis of the Living in Ireland data at a macro level, see Nolan, B., Whelan, C.T. and Williams, J. (1998).

number of Census variables which can be regarded even as indirect measures of poverty (e.g. percentage in long term unemployment) is quite small, so many indices tend to include indicators which measure the number of people with a raised risk of poverty (e.g. single mothers, elderly people living alone, etc.). However, there are areas in every city with high percentages of elderly people who happen to be quite affluent. So, although elderly people may in general be more at risk of being poor, a high percentage of elderly people is not necessarily a reliable indication that a particular area experiences poverty. Excluding the more indirect indicators of poverty is probably advisable, but a more fundamental problem is that different composite indices produce different results depending upon which constituent indicators happen to be included or excluded.

The method used to combine the various indicators into a composite index provides a second source of subjectivity. Multivariate indices require you to add together indicators that have different units, so you need to decide how much weight should be allocated to, say, each unemployed person relative to a household without a car when calculating the index. There is no generally accepted theory to provide guidelines, so in practice the weights are usually determined either a) arbitrarily (e.g. giving each indicator equal weight) or b) by some empirical means (e.g. principal components analysis). Decisions regarding the amount of weight to allocate to each indicator will clearly influence the value of the calculated index, and therefore which areas are identified as being in greatest need.

In many instances, principal components analysis or factor analysis is used to determine the weights. Using a complex multivariate statistical technique creates an illusion of scientific exactitude, but all it really does is disguise how arbitrary the method used to assign weights actually is. Both techniques, for example, tend to allocate weights to each indicator according to the extent to which it is correlated with (i.e. has a similar spatial pattern to) other indicators. Thus, if you have a number of indicators measuring more or less the same thing, they will be assigned higher weights and (because there are several of them) they will also be a double counted, irrespective of whether the dimension they measure is or is not a good indicator of poverty.

A further problem with using factor analysis or principal components analysis is that if you were to repeat the study for a later time period using exactly the same indicators, the weights assigned to each indicator would be different. This, coupled with the way in which these techniques calculate the scores for each area, means that the composite index for a given area could not be compared with the composite index for the same area at a previous point in time to gauge whether the situation had improved or disimproved. At best you might be able to say how a particular area was 'performing' relative to other areas (although even this would be open to question given the change in the weights), but it would not provide a very sound basis for evaluating whether a particular policy was cost effective.

The concept of a single poverty index which could be used to compare different areas at a particular point in time or to monitor the changes in a selected area over time is very attractive. However, in the absence of income data which could be used to calculate an indicator that everyone would understand (e.g. mean disposable income) then it is necessary to resort to surrogates. If these surrogates indicators are

to be combined into a single composite index, then they should ideally be combined in a consistent manner which is based upon an agreed theory of the relative importance of each indicator as a measure of poverty. Complex multivariate statistical techniques provide no substitute for such a theory. All they do is produce complicated, yet arbitrary, measures of nothing in particular.

In the absence of income data which could be used to measure poverty directly, I would suggest that it would be better to identify a number of simple indicators of poverty-related problems (e.g. unemployment, poor housing, educational deprivation, etc.). These indicators should be analysed separately, rather than combined into a composite index. They should be calculated in a way which everyone understands (e.g. simple percentages) and should also be designed to facilitate direct comparison not only between areas at present, but also facilitate direct comparison for different time periods.

Keeping things simple would also have advantages for determining policy. Even if an area is identified as severely deprived using a composite index (and even if it was assumed that other indices would also identify it as severely deprived), composite indices provide no indication as to how the problems might be tackled. However, if simple indicators are used, then it becomes much simpler to identify what sort of strategies might be required. For example, if an area has high unemployment, then the policy responses should be geared towards creating more jobs; if the area has a shortage of housing, then the response should be geared towards creating more houses; and so on. This is not to say that the policy responses should be either 'top-down' or uncoordinated. If an area has both an employment and a housing problem, then the most effective policy response might be a locally coordinated project to build houses using local labour. The main point is that the policy response should be geared towards tackling the identified problems, and not simply be a question of redirecting resources to an area to be used in a undefined manner because it scores high on a composite deprivation index.

The Spatial Framework

Attempts to identify areas of 'poverty', 'need', 'deprivation' etc. normally involve dividing the total study area (e.g. Ireland) into a number of smaller areas (e.g. counties) for which indicators or indices may be calculated. Those with the highest or lowest indicator values (as the case may be) are identified as being in greatest need. However, the choice of spatial subdivisions can have a major bearing upon the results.

The general rule of thumb is to calculate indicators for as small areas as possible, whilst remaining aware of the fact that some bizarre results might arise if the areas are too small because of what is often referred to as the small numbers problem. The smallest areas in Ireland for which census data are routinely published are referred to as DEDs (District Electoral Divisions) or wards with the County Boroughs (i.e. cities).[19] These are frequently used for geographical studies of poverty and related phenomena, but they introduce a number of problems which can easily be overlooked.

---

[19] The term DED will be used throughout the remainder of this paper to include urban wards, unless explicitly stated otherwise.

One limitation with DEDs is that they are essentially artificial areas. All political and administrative boundaries are to some extent artificial, but the boundaries between the larger areas generally correspond to differences in administration. For example, when you cross over a county boundary from one county to another, you move from the area administered by one county council to the area administered by another, possibly using different procedures which may have an impact upon poverty. People also identify with counties (and numerous other territorial divisions – e.g. parishes). However, DEDs serve no administrative purpose, and very few people could even tell you which DED they live in. This, in itself, is not a major problem – most alternative small areas would be equally artificial - provided that DED boundaries do not acquire a significance that is not justified. It would be foolhardy to deny an area which is obviously deprived from receiving special assistance simply because it fell on the wrong side of a DED boundary. There will invariably be situations where adjustments are required to whatever small area boundaries are used to get a justifiable outcome.

DEDs suffer from more serious deficiencies. One is that they vary considerably in terms of size (as measured by the population living within their boundaries) – in fact, the largest DED is over 1,000 times larger than the smallest. This in turn can result in a very misleading impression when the objective is to rank areas according to their score on a deprivation indicator. The reason for this is that the larger DEDs generally tend to be socially heterogeneous, therefore if they include deprived areas the deprived areas may be hidden by being included with more affluent areas. For example, suppose a DED contained two areas – one with a unemployment rate of 25 per cent and the other an unemployment rate of 5 per cent. If the two areas were the same size, the overall unemployment rate would be 15 per cent. Large DEDs will tend to have indicator values close to the national average (and therefore be excluded for special attention if, for example, targeting is restricted to the 10 per cent most deprived areas). Small DEDs, in contrast, tend to be less socially homogeneous and are therefore more likely to have low deprivation scores or high deprivation scores. Those with the low deprivation scores have a higher chance of being designated than larger DEDs which possibly contain many more households in need.

In Ireland the smaller DEDs tend to be found in areas of lower population density (i.e. rural areas). The wards in the larger cities tend to be quite large in terms of population, but they also tend to be socially homogenous, so the deprived areas in the cities tend to be captured in the indicators (see below). The areas that are especially disadvantaged by this particular problem are the medium sized towns, many of which fall within a single DED which scores close to the national average, but which contain clearly identifiable pockets of severe deprivation. These areas would have little chance of being designated for special attention if designation is dependent upon deprivation scores calculated for DEDs.

A somewhat similar type of problem, although one having a different impact, arises because comparisons are generally made between settlements of different sizes. Areas with small settlements (e.g. rural areas, villages) tend to be socially mixed, whereas large cities are socially segregated (i.e. large areas within cities are inhabited by people from similar social backgrounds, whilst people from other backgrounds live in different areas). DEDs within the larger settlements are more likely to have either very low or very high scores on deprivation indicators, whilst

DEDs in the more socially mixed areas are more likely to have deprivation scores close to the national average, and therefore be less likely to be designated. The extent of deprivation in a rural area may therefore be overlooked because of this particular problem, although the fact that many rural DEDs also tend to have small populations may also increase their chances of being designated (for reasons explained in the previous paragraph).

It is important to remember that the focus for any geographical analysis of poverty should always be people, not areas. It is very easy to fall into the trap of thinking of areas as 'having problems', rather than simply being places where people experience problems. Mapping poverty using traditional choropleth techniques can inadvertently place too much emphasis upon areas rather than people. Figure 2 shows male unemployment in Ireland mapped by DED. One's eye tends to be drawn to the large areas of high unemployment in the west. This would appear to create a strong argument for new jobs to be diverted to Donegal, Mayo and Galway. The other map in Figure 2 shows the same data, but displays it as a cartogram in which each area is represented by a circle, the size of which is in proportion to the number of people living in the area.[20] The visual impression is quite different. Many of the areas which stood out in the previous map, especially those in Mayo and Galway, have virtually disappeared because, although they have high unemployment rates, relatively few people live there. The western suburbs of Dublin in contrast, which were previously invisible, emerge as major concentrations of unemployment. If the objective is to help people in need, as opposed to areas in need, then there is an argument for prioritising these areas for new job creation.

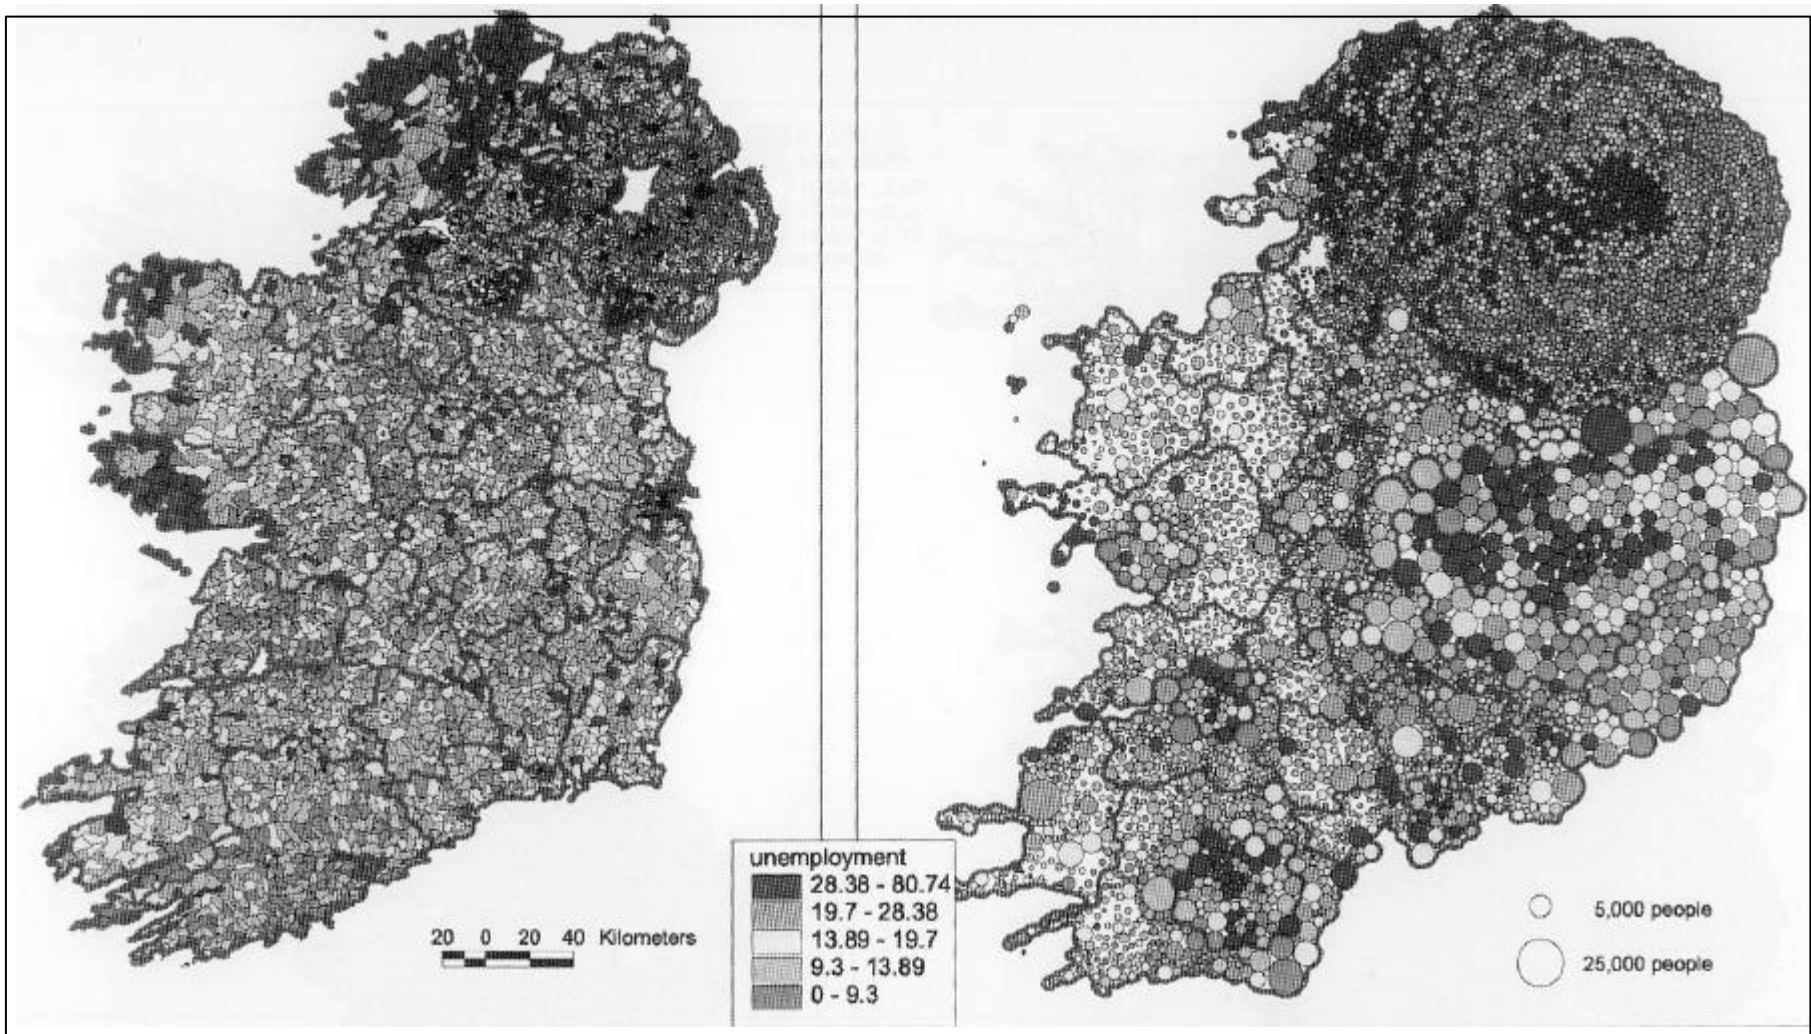[20] Source: Cook et al. (2000).

Figure 2. Male Unemployment by DED

Most of the problems discussed here arise because we divide the country into areas for the purpose of calculating indicators. However, it is possible to explore the geography of poverty without the use of essentially arbitrary areas. Image a circle with a specified radius (e.g. 20 miles). If one had information on each individual in the state geocoded as a point corresponding to their home address, one could count what percentage of the people living within the circle were unemployed or, if one had income data, one could calculate the mean disposable income of people living in the circle. These calculations could then be repeated for a second circle of similar size drawn around a second point displaced a few miles from the centre of the first point. By repeating this for a large number of circles, one could build up poverty indicator measures for a large number of points. These indicator values could then be 'contoured' to produce a map with 'hills' representing high levels of deprivation and 'valleys' representing low levels of deprivation. Additional resources could then be channelled to the 'hills' in the map.

One advantage of using this approach is that by using different sized circles, one could explore the spatial dimensions of different forms of deprivation. Some forms of deprivation are localised and require localised solutions; others affect larger areas and need to be tackled at a different spatial scale. However, this approach could only be implemented if the information on each individual was geocoded – i.e. if each person's address was recorded as a grid reference. This is technically feasible: Ordnance Survey Ireland and An Post have developed a computerised geocoding system which in theory allows each address in the state to be translated into a grid reference. However, a decision would need to be taken at governmental level to encourage all data collecting agencies to use this or a similar system to geocode their data. Funding would need to be provided to fund the software.

Steps would also need to be taken to guarantee the confidentiality of individuals who could easily be identified by anyone having access to their personal details containing a geocode. Whilst it is difficult to envisage the government grasping so many nettles simply to facilitate research on poverty, there are many other reasons why geocoding should be encouraged.

Policy Implications

If areas of poverty or deprivation could be objectively identified, then the worst affected areas could be designated to receive additional funding to help counteract their disadvantages. However, questions obviously arise as to how this funding might most usefully be utilised. I would like to finish with a few thoughts on possible policy strategies.

The first point I would make is any strategy based upon positive territorial discrimination must be seen as a supplement to, rather than a replacement for, the normal universal mechanisms for providing assistance to people in greatest need (e.g. welfare benefits, medical cards, etc.). Typically only a minority of deprived people live in deprived areas, and only a minority of people in a deprived area are actually deprived. Identifying areas of deprivation is therefore not a very efficient method of ensuring those most in need actually receive assistance. Assistance, especially assistance to cope with immediate problems, should therefore be primarily

directed at identified individuals in need rather than at designated areas. An unemployed person is still unemployed, and in need of assistance, whether they live in an affluent area or a deprived area.

Nevertheless, people's life chances are influenced by where they live. The chances of being educationally deprived, unemployed, low paid, in poor housing, in fear of violence, or in poor health are all obviously much greater in some areas than in others. Such areas should be identified and designated for special treatment to counteract their *structural* disadvantages – i.e. to counteract the processes which result in the area being disadvantaged.

Different strategies will be required depending upon the type of problem (or problems) that need to be addressed. Also, the specific causes of the problems may often be different in different areas (e.g. high unemployment in one area may arise because of a high number of school leavers, and in another area might arise because of a decline in a traditional industry, such as agriculture). Different local areas may therefore require different strategies to alleviate their problems. I would suggest that local people usually have a good idea of what their problems are and how they should be tackled, even though they may lack expertise in some areas or may benefit from advice from outside. The impetus for change should therefore, where possible, come from within the area (i.e. be 'bottom up').

One problem with many local initiatives is that there is sometimes a lack of imagination and a tendency to replicate what has worked in other others. It certainly makes sense to benefit from the experience of others, but there is a limit to what can be achieved by simply replicating successful ventures elsewhere. For example, every area now seems to have its own local tourist initiative, but there is a finite number of potential tourists, so whilst tourism may provide a solution to some of the economic problems for areas with an immediate appeal for tourists, the same success is unlikely to be replicated in other less naturally endowed areas. Such areas may need to capitalise on other resources. Local people are again in the best position to come up with new ideas. The point I wish to make is that they should be encouraged to be imaginative. Local area schemes should be regarded as an opportunity to experiment, without too many constraints being placed upon the projects to deliver short-term results.

Government clearly has an important role to play.[21] I would see government fulfilling at least three essential roles. First, it should try to ensure that resources are allocated to areas according to their need, rather than simply to those who shout loudest. This would require the development of indicators along the lines discussed in the previous sections of the paper. Second, given that projects are supported from public funds, there is a requirement to ensure that these funds are not being misappropriated. However, this needs to be handled sensitively. There is a very fine line between monitoring how resources are utilised and such monitoring becoming overly bureaucratic or used as a top-down management tool (e.g. by withholding funding from projects that do not satisfy specific criteria). Thirdly, government has an

---

[21] The term 'government' is used loosely here to refer to central government, local government and European Community bodies and agencies of all types, in contradistinction to local community groups, whether formal or informal.

important advisory role to play, drawing for example on the experiences of experimental projects which may or may not have succeeded.

It is important to be aware of the importance of spatial scale. Schemes to tackle certain types of problems may require a very localised strategy (e.g. at the level of a village, or a single housing estate in a county town). Other schemes will require a broader based approach, maybe at the level of a county or several counties. The funding mechanisms should be flexible enough to accommodate schemes at a variety of scales.

Finally, it must be recognised that there is a limit to what can be achieved at local level. Just because a particular problem may be concentrated in a particular local area is no indication that the causes of the problem, and therefore its possible solutions, originate from within the area. Many localised problems are simply the localised manifestation of processes that operate at a national or global level. Whilst area-based strategies may be able to alleviate some of the symptoms of these problems, there is very little that they can do about the causes.

Recommendations

The principal recommendations arising from this review may be summarised as follows:

- There is a need for objective studies of the spatial distribution of poverty.
- These studies should use measures which facilitate not only comparisons between areas at a given point in time, but also comparisons within areas between different time periods.
- Indicators should be measured in a way which permits temporal comparisons.
- Poverty should ideally be measured by an income measure
- If surrogate measures must be used, then they should be analysed separately. At least that way you have some way of knowing what you are looking at.
- DEDs are not very suitable for small area analysis. They should be replaced, as a minimal requirement, by a set of smaller areas more uniform in size for census and other data collection purposes.
- Geocoding would permit indicators to be calculated for areas of any given size or shape. It would also permit more detailed analyses using 'moving circles' of different sizes. Government agencies should be encouraged to adopt geocoding for record keeping purposes.
- Area-based schemes should, where possible, capitalise on local initiative and expertise to counteract the structural causes of poverty and deprivation in the worst affected areas.
- Different problems require different solutions in different areas. Government has an important role to play in fostering a diversity in approach.

References

Cook, S., Poole, M.A. Pringle, D.G. and Moore, A.J. 2000. *Comparative Spatial Deprivation in Ireland: A Cross-Border Analysis*. Oak Tree Press, Dublin.

Haase, T. 1995. *The Designation of Disadvantaged Areas in the Local Development Programme*. Report to Area Development Management, Ltd.

Haase, T. 1999. 'Affluence and deprivation: a spatial analysis based upon the 1991 Census of Population' in Pringle, D.G., Walsh, J. and Hennessy, M. eds. *Poor People, Poor Places. A Geography of Poverty and Deprivation in Ireland.* Oak tree press, Dublin.

Nolan, B., Whelan, C.T. and Williams, J. 1998. *Where Are Poor Households Found? The Spatial Distribution of Poverty and Deprivation in Ireland.* Oak Tree Press / Combat Poverty Agency, Dublin.

Pringle, D.G. 1998. Hypothesized Foetal And Early Life Influences On Adult Heart Disease Mortality: An Ecological Analysis Of Data For The Republic Of Ireland. *Social Science and Medicine*, 46(6), 683-693.

Small Area Health Research Unit (SAHRU) 1997. *A National Deprivation Index For Health And Health Services Research*. Dept. of Community Health and General Practice, Trinity College Dublin.