High resolution colour mapping using modern technology has allowed us to explore the breadth of census data available to the contemporary social scientist. An examination of methods used to create cartograms, which minimise visual bias involves considering how densities and area boundaries should be mapped. The Modifiable Area Unit Problem is discussed and the advantages of using three-colour or trivariate mapping are outlined. Further modifications to traditional census mapping such as the mapping of change and flow are addressed.

Map design for census mapping

D. Dorling

Department of Geography, University of Newcastle upon Tyne, Newcastle upon Tyne NE1 7RU

For maps of larger scale, an artistic objective might well lessen our insistence on a strict geometric framework for maps and make room for the greater use of mental constructs of social, cultural, and economic space . . . Such maps might well be considered the cartographic equivalent of 'mild' surrealistic art.

Arthur Robinson (1989 p.97)

INTRODUCTION

This paper aims to give a flavour of new cartographic techniques which can be used to study census data. Techniques used to study differences between many small places which often require colour are the paper's main focus. What makes the census particularly valuable for social science is not the breadth or depth of the questions that it asks (as they are few and shallow), but the great spatial detail that is provided - showing how each neighbourhood, each block of streets, each hamlet, differs socially from its neighbours (for every place in the country simultaneously). The census also acts as the base for almost all other mapping of British social, economic, political, housing and medical statistics as it tells us where households, jobs, voters, dwellings and people are. Whilst we cannot have information on every individual, due to confidentially constraints, we are given many statistics for small groups of people aggregated into geographical areas. The question being addressed here is "how can we map the detail this data contains but in its raw form conceals?"

The 1851 British census was the first to be mapped in any detail (Petermann, 1852), although American censuses were mapped before this and ancient maps can be found of the population in, for example, China, *Figure 1* shows a dot map of England and Wales showing the distribution of the people in 1851. Graphical depiction of statistics was not in favour over much of the subsequent century (Beniger and Robyn, 1978). The next British census to be mapped nationally and in detail was that of 1961 - by hand drawn choropleth maps using the local authority boundaries of the day (Hunt, 1968). Computers were first used to map the 1971 census. *People in Britain: a Census Atlas* (CRU, 1981) depicted characteristics of the population by colouring kilometre grid squares across the whole of Great Britain to show the detailed spatial structure of our society. In 130 years, however, the fundamental look of maps of the population had changed little. Since then, with two decades of development in computer graphics, a great deal has changed.

PLOTTING POINTS

The finest level of output of the 1991 census spatially is the 'enumeration district' in England and Wales and the 'output area' in Scotland. Each enumeration district contained, on average, four hundred people and consisted of just a few streets or a single block of flats in towns and cities, part of a village or a single hamlet in the countryside. Output areas are even smaller. There were over 131,000 enumeration districts in 1991 in England and Wales and 26,000 output areas in Scotland. We have been given information about very many small groups of people and all their geographical neighbours.

For each enumeration district a national grid reference to an accuracy of one hundred metres is given in what are called the *small area statistics*. This identifies the centre of population of the area. The points were chosen by hand and in the first release of the statistics several thousand were found to have been misplaced (Atkins and Dorling, 1993). *Figure 2* shows the corrected spatial distribution of these collecting units and through them the national population distribution. All the figures, apart from *Figure 1*, were produced on a home microcomputer using some very basic programming. The practicalities of census mapping are now a lot simpler than they were in 1851 (although disc space was not a problem then!); now it is the theory of how better to represent the population which is most challenging.

Along with the 1991 census data, digital boundaries of enumeration districts have been produced, in theory delimiting precisely where 'on the ground' people were counted (although the academic purchase of these are still

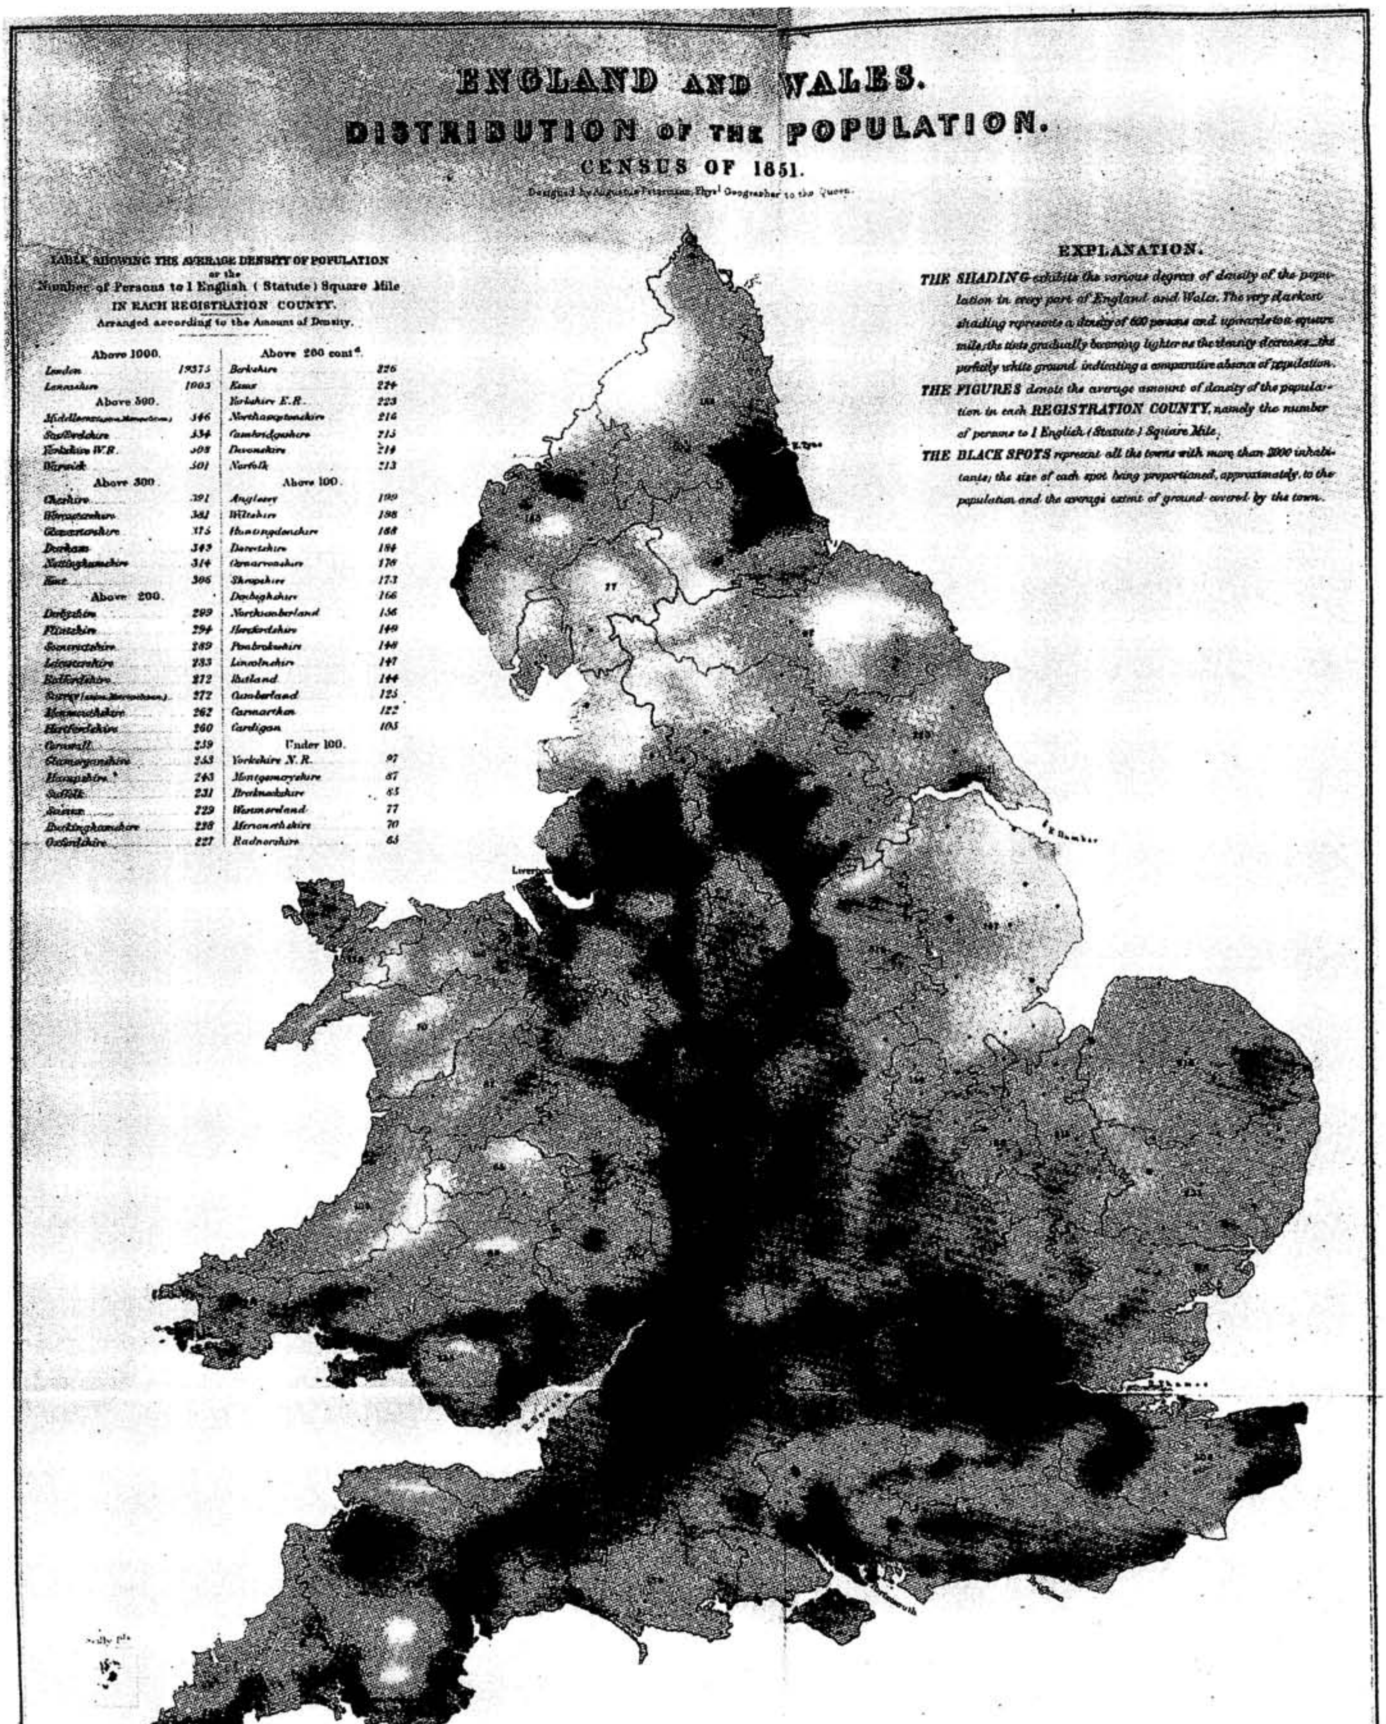

Figure 1. England and Wales, Distribution of the Population, Census of 1851.

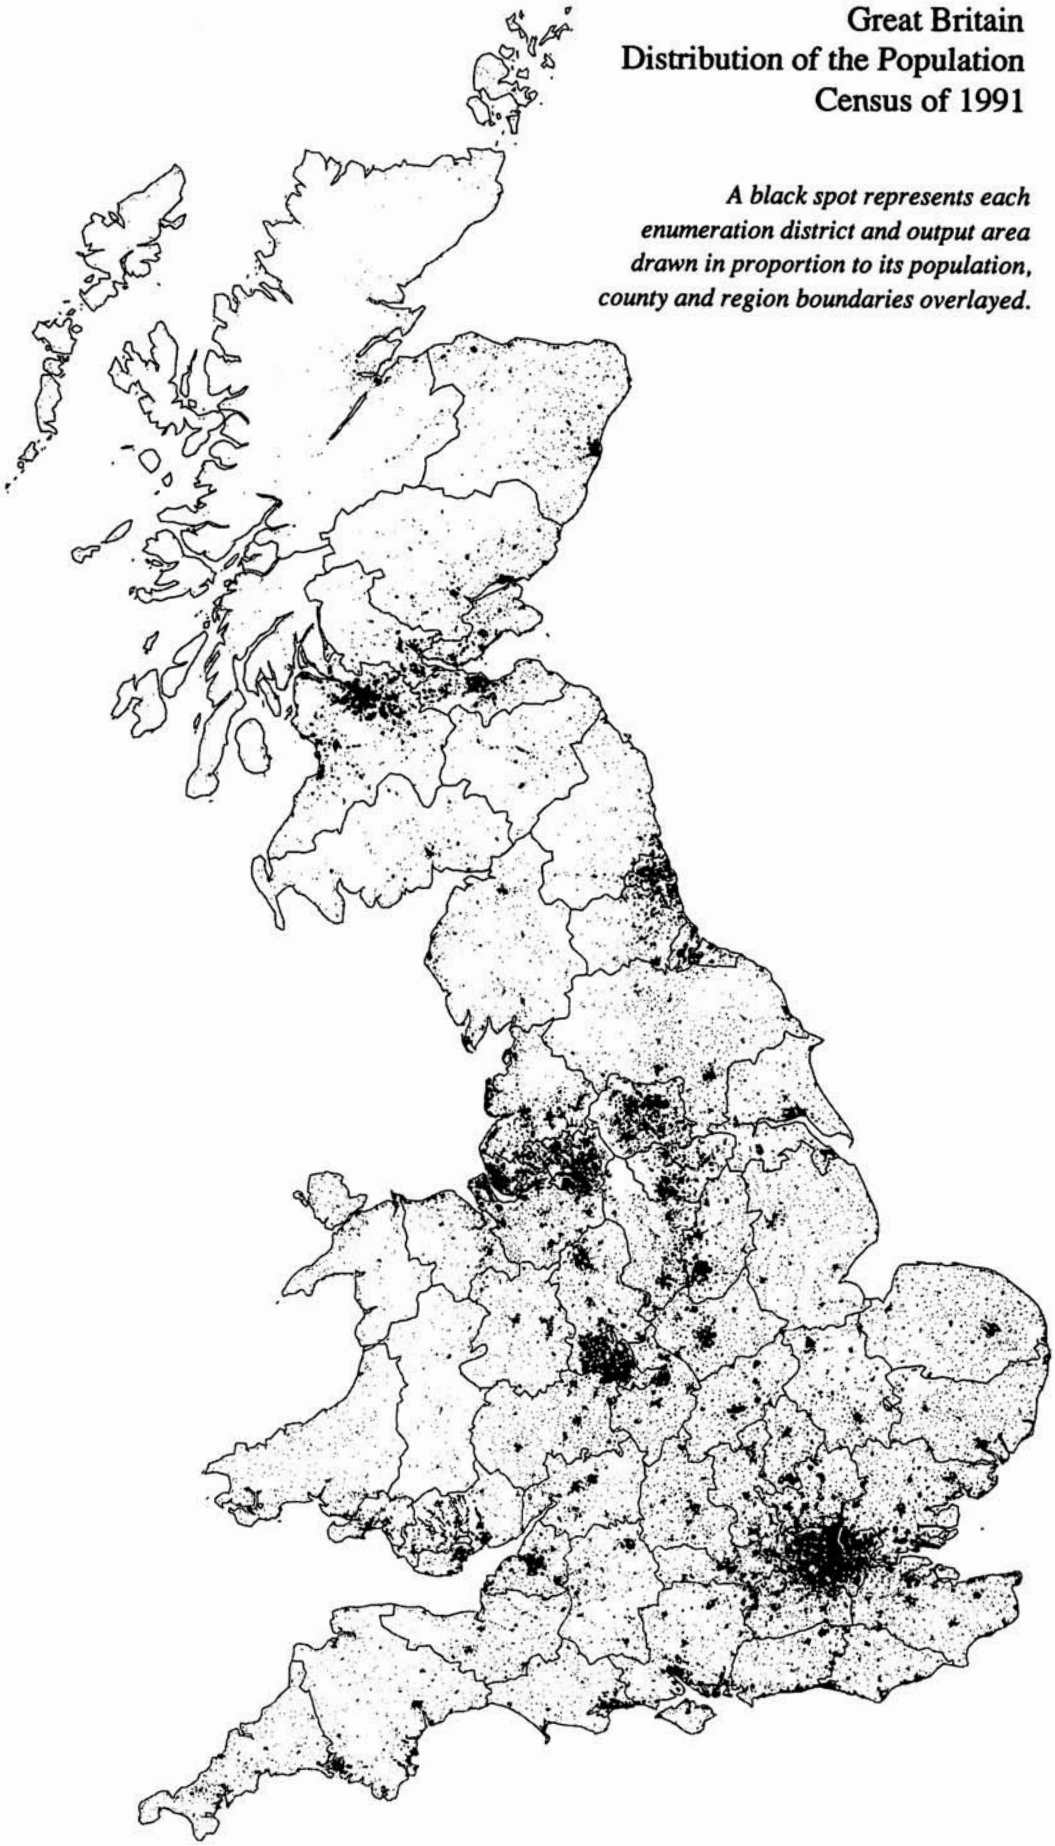

Figure 2. Great Britain, distribution of the population, census of 1991.

due for release at the end of 1993). Enumeration districts range in size from less than one thousandth of a square kilometre (the base of a high rise block) to over one hundred square kilometres (encompassing empty moorland in the most remote areas). It would not be surprising to find that the largest is more than a million times the area of the smallest whilst still containing fewer people. Mapping with these boundaries will no doubt take place, but could produce some very misleading pictures because these areas encompass all the empty land as well as that occupied by people. Ian Bracken and David Martin have suggested some methods using Geographical Information Systems to avoid this problem (Martin, 1991). Here I concentrate on alternative cartographic solutions.

Dasymetric mapping – shading only inhabited parts of a map – is seen as one possible way forward (Langford *et al.*, 1990). One solution is to take the dasymetric mapping of the 1851 census to its digital limits and attempt to draw a point to represent every household. Within enumeration districts we have a good idea of where the households are because a postcode is now linked to every census form and we know the rough location of the addresses allocated to each postcode. *Figure 3* shows an inset of central London

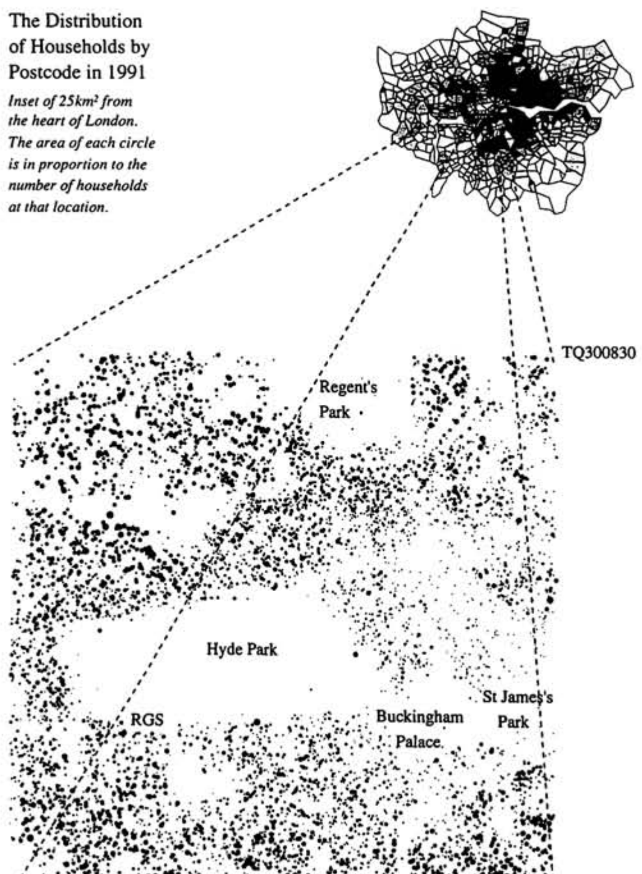

Figure 3. The distribution of households by postcode in 1991.

with a circle drawn at the location of each postcode, its area in proportion to the number of households living there. Unfortunately, as we portray physical reality more accurately, we are left with less and less space in which to show social reality. Even in the most densely populated part of Britain what dasymetric mapping shows most clearly are those areas which it leaves blank. The parks, rivers and roads at the centre of the capital are what stand out clearly on a dasymetric map of any variable for a city like London, while outside the cities the isolated places are most prominent, places which usually contain least people and for which averages and proportions are least meaningful.

How might we go about visualizing more fairly the characteristics of people in so many areas, which vary so much in physical size? *Figure 4* shows a grid drawn over a

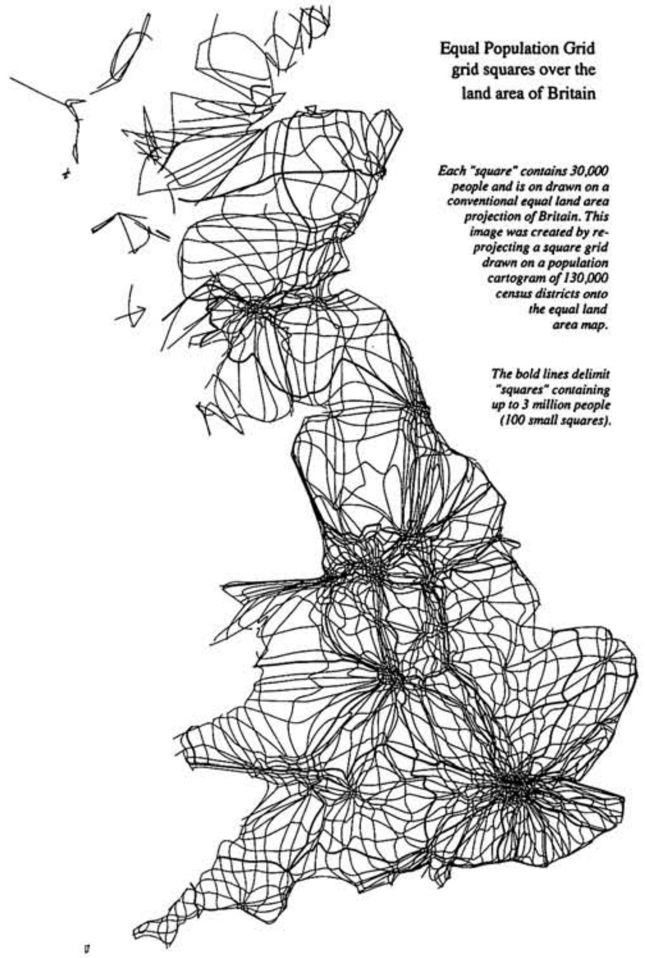

Figure 4. Equal population grid squares over the land area of Britain.

map of Britain. The grid is not uniform, but is violently stretched and twisted. It is made up of almost two thousand squares each of which (on the mainland) contains 30,000 people. There are many such squares in the large cities, but few in the more rural areas. The picture is more complex than this, however; for in East Anglia, for example, the effect of Norwich can be seen, like a weight pulling the fabric of some giant net inwards. Each square contains the same number of people, yet some cover great swathes of land while others are barely visible. Perhaps it would be more sensible to pull all the lines straight – so as to form a rectangular grid – ensuring that each square would be of equal size so that their populations are more fairly represented than on a traditional map. This would create an *equal population cartogram* (Tobler 1973) which is discussed later.

DRAWING BOUNDARIES

For the 1981 census, digital boundaries were available only for wards, of which there were 9,289 in England and Wales and into which enumeration districts were nested. The electoral wards varied in population from 498 (Lower Swaledale) to 41,502 (Birmingham, Weoley), and in land area from 19 hectares (Skipton Central) to 44,789 hectares (Upper North Tyne). They exhibit a skewed distribution in which the majority are small in area where people huddle together on the land. On a conventional national grid projection only a small proportion of the wards are visible. *Figure 5a* shows all wards' boundaries for the whole

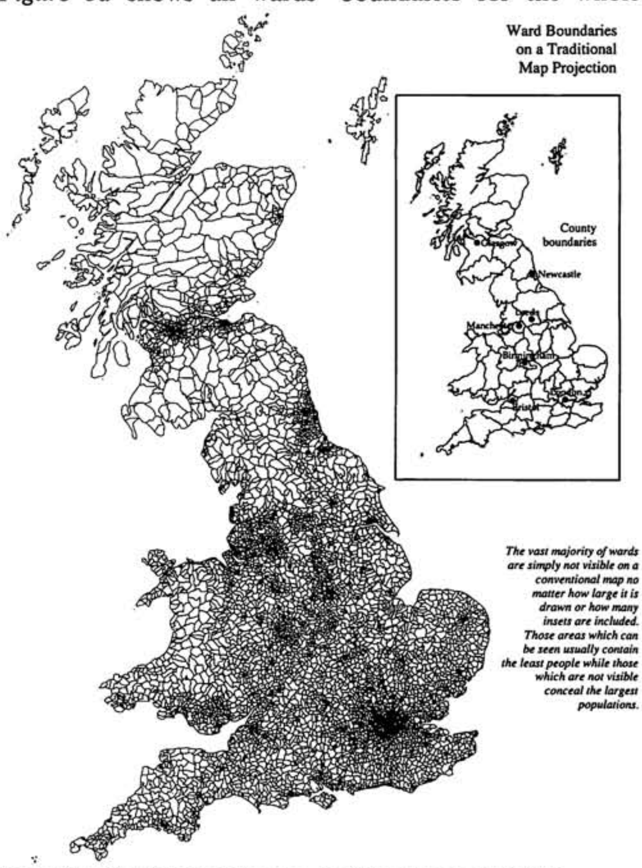

Figure 5a. Ward boundaries on a traditional map projection.

country (the inset in *Figure 3* focuses in on London). At almost every level of magnification, below one thousand fold, there are some wards which cannot be seen and others, usually containing the least people, which dominate the image. Simply because the census provides geographical information using one projection does not mean that that projection must be maintained, although almost all recent mapping does this. Similarly, we do not have to use the raw counts of the people that the census provides, but can transform these to use more meaningful scales.

Chi-squared mapping (Visvalingham, 1981) is a good example of how shading categories can be altered in an attempt to overcome the overemphasis traditional maps give to the characteristics of unpopulated areas. Traditional thematic mapping makes most use of ratios to characterise areas, for instance of the number of unemployed people divided by the total workforce. A signed chi-square statistic can be calculated in place of the more usual ratio for each area and areas are only shaded brightly when this statistic differs significantly from what would be expected. This avoids the basic problem with census statistics, that the most extreme ratios (and hence the areas which draw most attention on maps) are almost always found where there are the least people and hence where the statistics are most unreliable. Chi-squared mapping was used to map the 1971 census (CRU, 1980), and is a good palliative if mapping with an equal land area projection is a high priority.

The information in the census concerns not land but people and households. In visualizing these, a primary aim can be that each person and each household is given equal representation in the image. Current computer versatility allows this, so a principal challenge of visualizing the 1991 census is to achieve this. *Figure 5b* shows an equal population

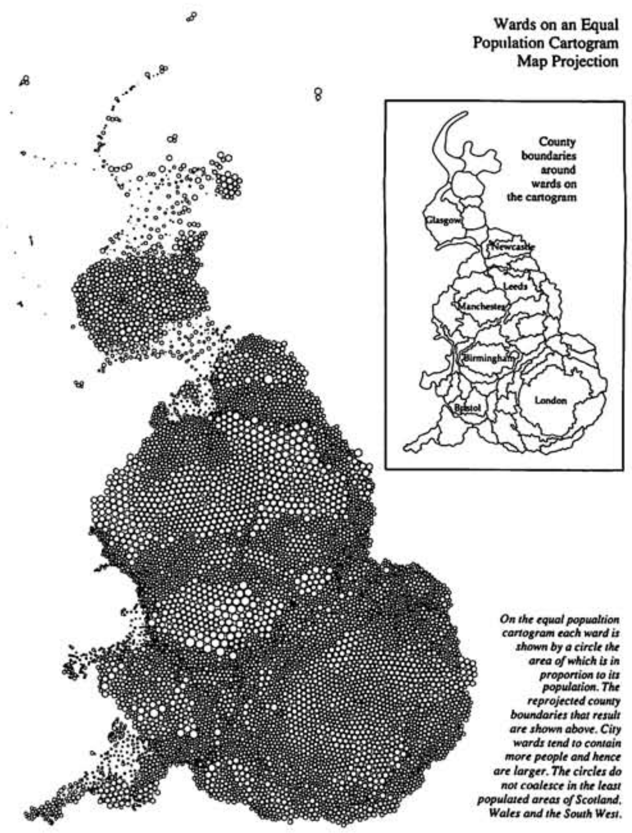

Figure 5b. Wards on an equal population cartogram map projection.

cartogram of the ten thousand wards of Britain based on the 1981 census. In effect, the lines in *Figure 4* have been straightened (for more details see Dorling, 1991). On the cartogram each ward is represented by a circle, its area in proportion to its population and located as close to its original geographical neighbours as possible. The cartogram represents one level of abstraction beyond the equal area choropleth map – one level closer to the maps of social landscape we might aim to see if we wish to look both inside cities and across the nation, simultaneously.

Many arguments can be made for why we should use cartograms to portray spatial distribution from the census. By trying to minimise visual bias, cartograms can be claimed to have advantages in census mapping, being more sensible statistically and more just socially. Their main disadvantage is that they are unfamiliar, but we do not learn from familiarity. Traditionally cartograms have been used in mapping election results and medical statistics (Hollingsworth, 1964; Howe 1970). This is because, theoretically and respectively, every elector's vote is equally important to the outcome of an election and we are concerned with the spread of disease among individuals, not over land. The British Labour party will never appear to have won on a traditional map of constituencies while incidents of disease always appear massively concentrated in cities on a pinmap. If you believe that every individual's circumstances are equally important in the spatial make-up of a society, then you should use area cartograms in mapping census data.

REPRESENTING PEOPLE

However carefully we define the value of a variable and assign a shading category to an area we will still be using arbitrary spatial boundaries for calculating that value in the first place. This dilemma is well known as the Modifiable Areal Unit Problem (Openshaw, 1982). For census mapping the Modifiable Areal Unit Problem can be handled in a number of ways. The simplest is merely to illustrate it by using multiple boundaries, and with cartographic computer animation it is possible to redraw images instantly using different boundaries, to see the effects of these choices (Dorling, 1992). A more sophisticated approach to this problem is to think more carefully about what areas are meaningful to visualize. A set could be specifically defined (e.g. Housing Market Areas), or 'fuzzy boundaries' could be used – for instance by employing kernel mapping methods (Brunsdon, 1991).

Arbitrary boundaries, however, have their greatest influence on the impression gained in their use for portraying statistics, not in their use for calculating them. If the whole of a region is shaded dark grey because levels of unemployment are particularly high in one of its towns, is our image accurately reflecting reality? Just because proportions are calculated for one area does not mean they have to be shown using a scaled down replica of the boundary of that area (i.e. by choropleth mapping). David Martin (1991), for instance, makes a strong case for almost never using choropleth mapping.

The severity of the Modifiable Area Unit Problem for conventional thematic mapping is due, as has already been touched on, to the fact that greatest emphasis is given to those places containing fewer people – where the arbitrary movement of boundaries can have the most influence on the values calculated. This is illustrated graphically below. In general, arbitrary boundaries can be seen to have much less influence on the impression gained from cartograms as opposed to traditional maps (Fotheringham (1989, p. 223) has claimed that visualization may provide a range of solutions to the problem). Our methods of visualization must be robust if one arbitrary choice or another is not to lead us to misinterpret what is important.

The 1991 census mislaid over a million people who are thought to live in Britain. This may appear disastrous, but our conventional mapping techniques are capable of hiding many more than this from our eyes. Conventional maps *lose* people because they are designed to show land and other features of traditional military significance such as hills, roads and woodland (Harley, 1989). People are concentrated in cities in such small areas of space that, when we map their distributions nationally and in detail, the characteristics of the majority are almost impossible to see. When we map at higher scales we convert millions of statistics into only a few dozen or a few hundred numbers and then draw a picture of them – a great deal of information is lost. Pictures can contain much more information than this. The problem is to structure the information spatially in such a way as to maximise the content that can be seen (while still retaining the geographical topology) – to clarify reality through the portrayal of detail (Tufte, 1990, p. 37).

In the cartograms shown here most wards have been placed adjacent to wards with which they share a common boundary, although occasionally this has not proved possible (thus they are 'noncontinuous area cartograms' (Olson, 1976)). Even when reproduced to the page size of this journal every ward is made visible in a single image. Here is a projection upon which it is possible to show the fortunes of every group of people at something akin to the neighbourhood level, nationally, without a particular bias against those groups who happen to live in the most densely populated areas.

Figure 6 shows two images which give an impression of where some of the people who were missed by the census might be. They show the proportion of residents who were 'imputed' to exist by the census authorities when nobody was found in a dwelling and it was thought by the enumerator not to be vacant. On the map, these people can be seen to be concentrated in cities; but the largest *areas* of underenumeration are on the west coast of Scotland. The cartogram shows that this is a false overall impression to gain of the spatial distribution of the location of imputed residents. Underlying the cartogram is the true distribution of the population so an indication of the actual relative scale of the problem in each area is given. More than a third of all those imputed nationally were in London. This can be estimated by comparing the sizes of the shaded areas which are in proportion to the populations.

MAPPING SOCIAL STRUCTURE

The main subject of interest in mapping the census, however, is the characteristics of those people who did complete their census forms. These characteristics, when in turn used to characterise places, are often quite complex variables to map. For instance, although any single individual might be assigned to a single social class out of a three-fold division of occupations, the place in which that individual lives could be made distinct by the combination of the relative numbers of people in each of those three groups who are also living there. The census allows us to measure how much mixing of different social groups there is and how that is changing. There are statistical techniques to reduce such numbers down to a single index of, say, segregation, but information on the nature of social divisions is lost when these techniques are applied. The cartographic solution that census mapping has produced is to use different mixtures of colour to represent different mixtures of people. The US Bureau of the Census (1970) pioneered automated bivariate colour mapping to show how two characteristics of the population were spatially related to one another. Colour, however, contains three primary components and so can, in theory, be used to show the relationships between three dependent variables.

Three-colour, or trivariate, mapping is the most sophisticated possibility when there is only a pinprick to shade. Three primary hues of varying intensity will produce a unique colour for every possible combination (while more than three colours do not). The three primary hues I choose to use are red, blue and yellow (mixing to purple, orange and green) as I find these most intuitive (as does Arnheim (1970, p. 30)). Sibert (1980, p. 214) prefers red, green and blue, while Bertin (1981, p. 163) recommends cyan, yellow and magenta. Whatever choices are made, when images use three-colour shading, the patterns formed are usually much more subtle than those found in univariate structures. It is often the case that the more complex a picture is, the more there is to learn from it. This process is not necessarily easy, but can be very rewarding in terms of gaining new insights.

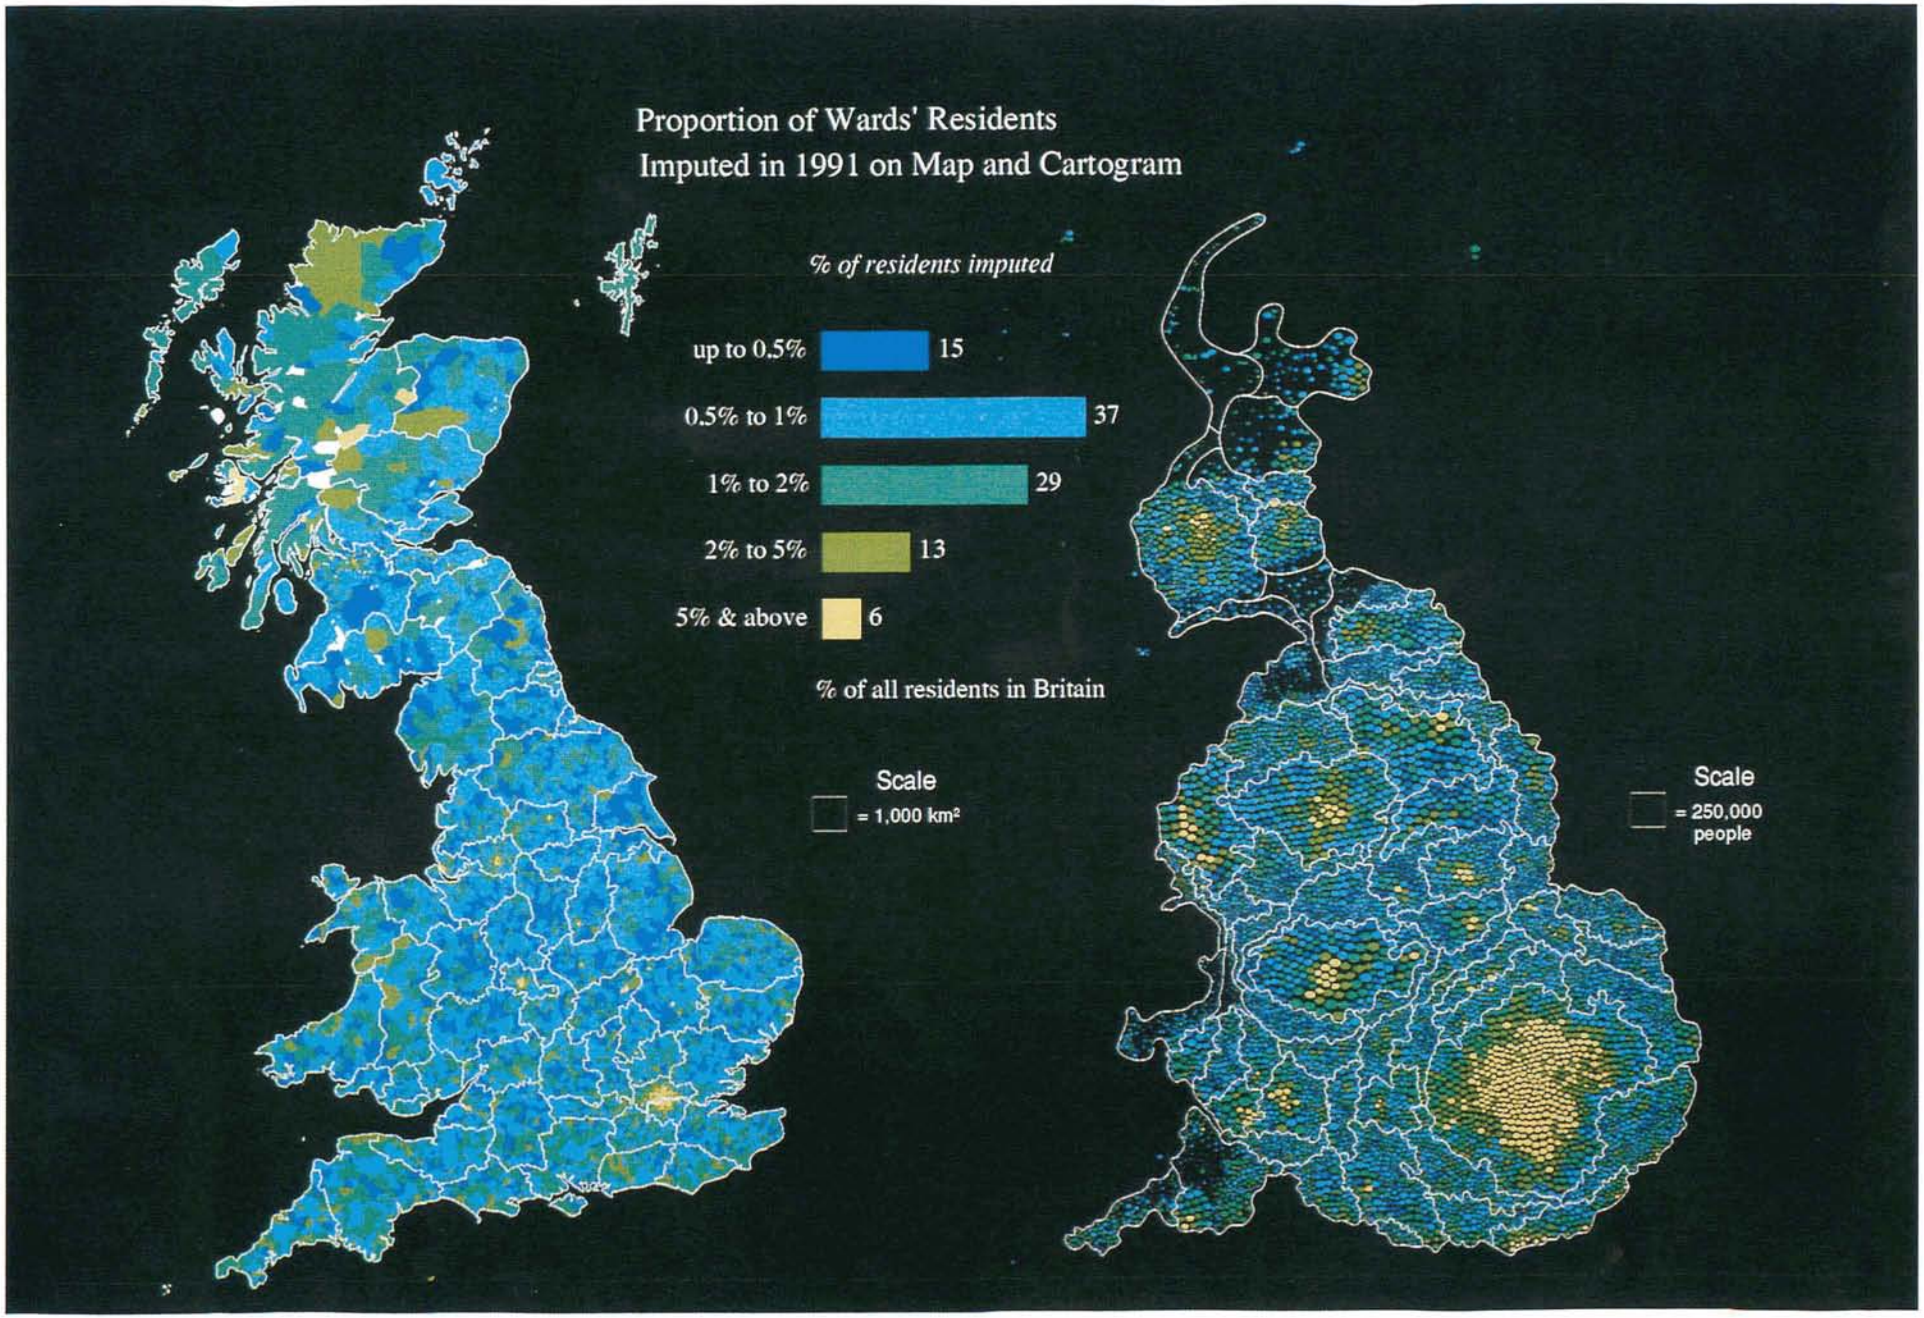

Figure 6. Proportion of wards' residents imputed in 1991 on map and cartogram.

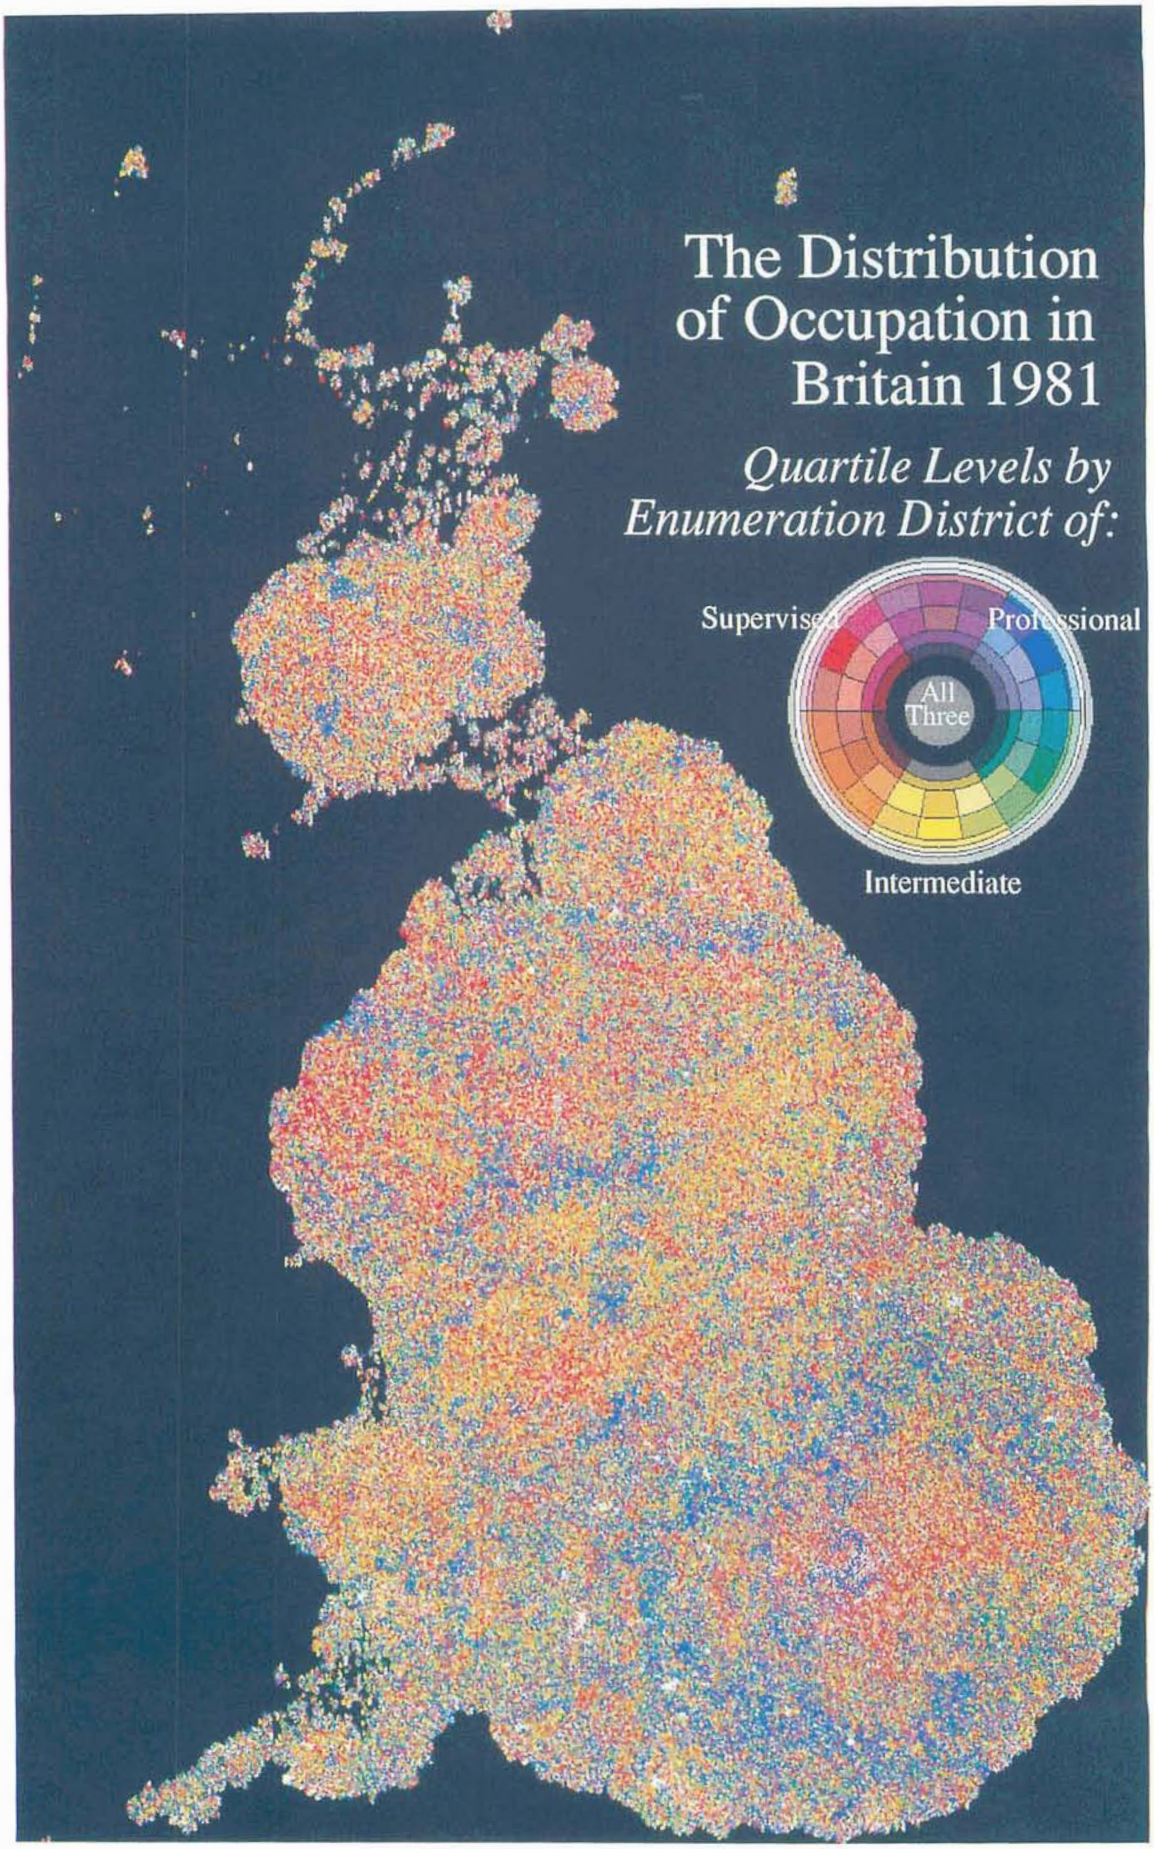

Figure 7. The distribution of occupation in Britain, 1981.

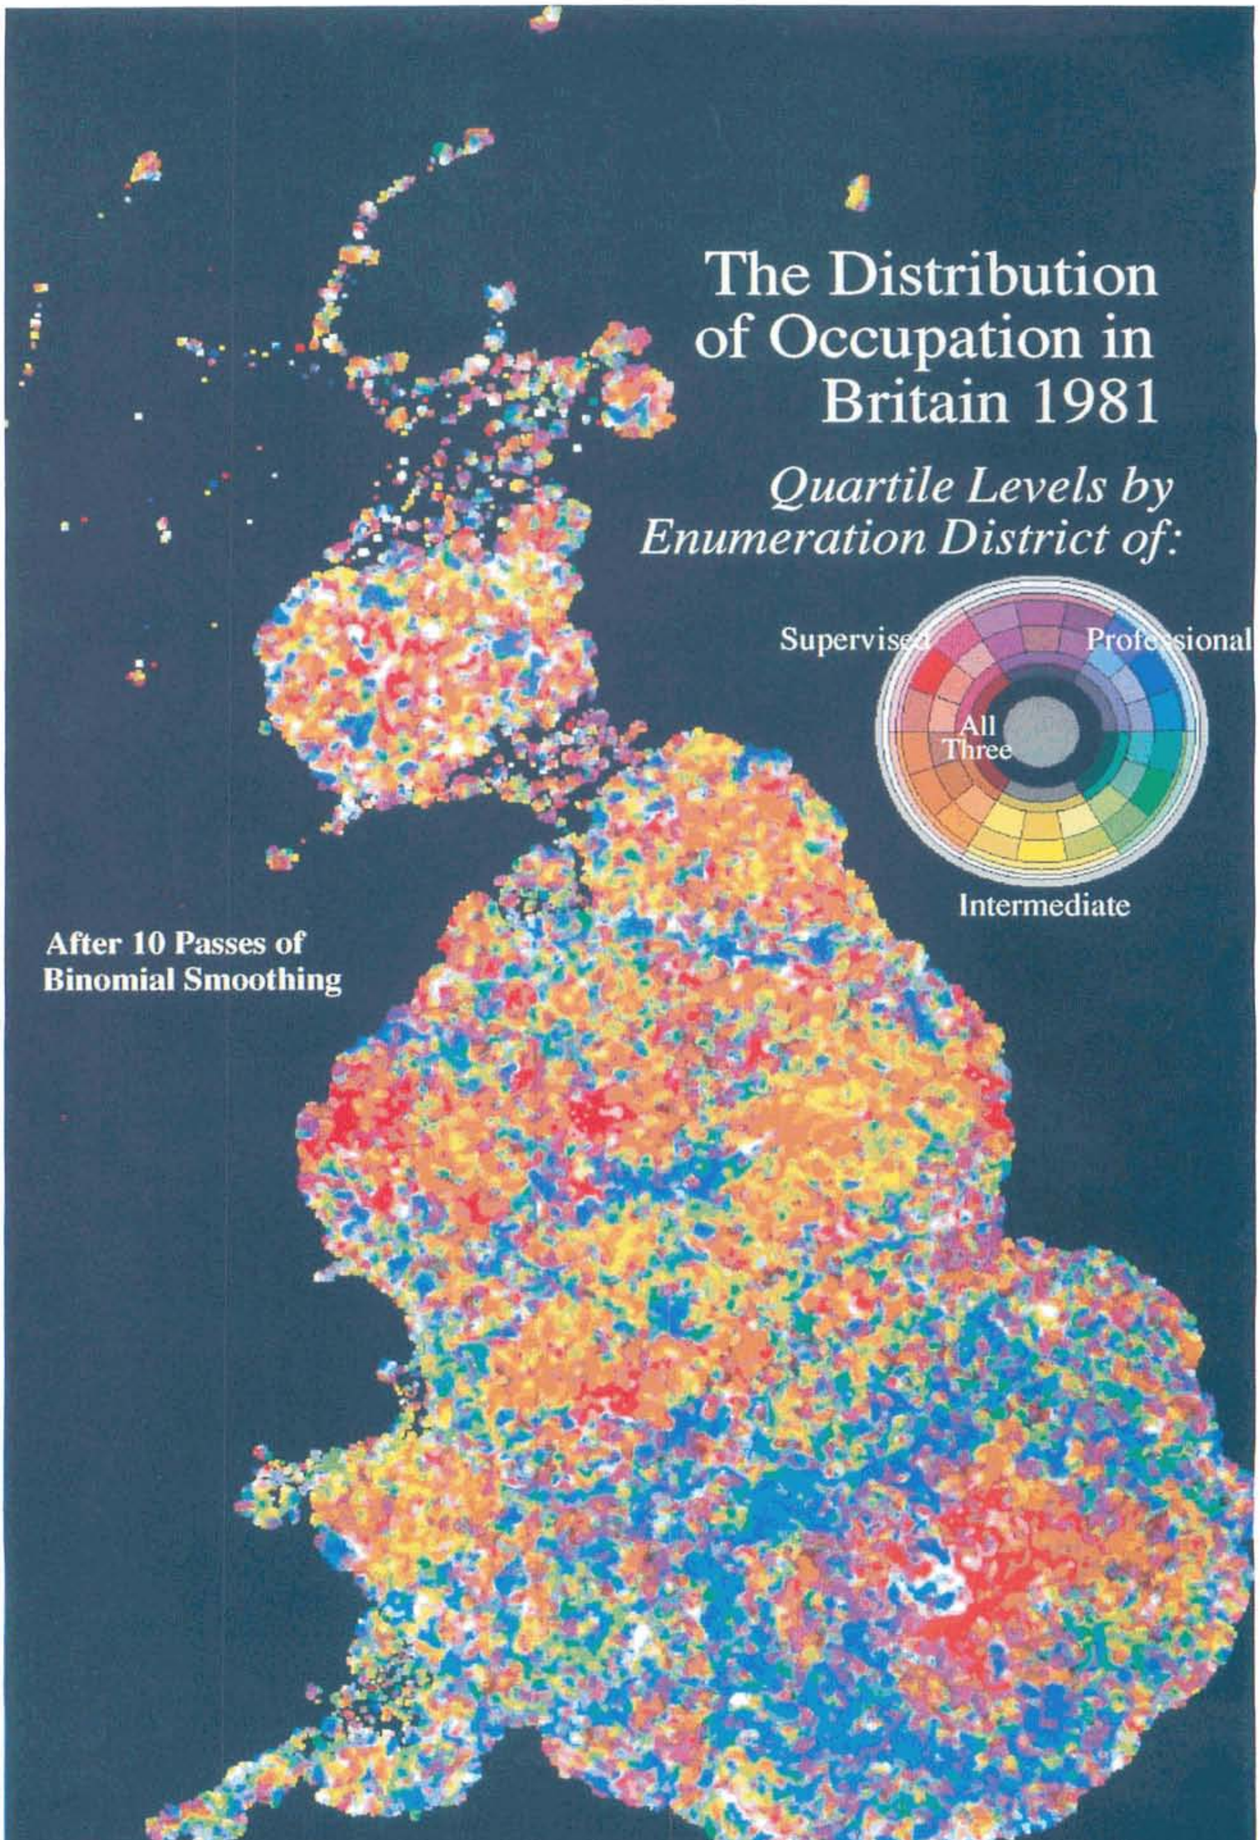

Figure 8. The distribution of occupation in Britain smoothed from EDs.

Figure 7 uses a population cartogram of the 129,000 1981 census enumeration districts where each district is represented by a circle coloured one of sixty four possible shades to show the detailed spatial distribution of three occupational social classes in Britain. Clusters of enumeration districts form where the occupation of 'heads' of households is predominantly 'professional' (blue), 'intermediate' (yellow) or 'supervised' (red). Where there is a concentration of professional workers, for instance around Greater London, the area will appear bright blue. Where there are high proportions of people in both intermediate and supervised occupations (but not in professional occupations), for instance around central Birmingham, areas of orange emerge. Superficially, there is not a great deal of difference between this kind of image and a CAT (Computer-Aided Tomography) scan, with the latter showing a slice through the human brain and the former, a slice through society. One arguable difference is that a great deal of training is needed to interpret a map of the brain, with which most of us are unfamiliar, whereas a map of society can be much more easily, understood superficially (as we, ourselves, constitute society) and so can successfully show more complex structures.

Because the three distributions of occupations can be seen to spatially 'repel' each other (people in different classes of job tend not to live in the same streets), the primary colours chosen, of red, blue and yellow, dominate this image. Central London, Birmingham, Liverpool and Manchester have distinctly high proportions of people in 'supervised' occupations, although a 'snake' of professionals winds its way, north to south, through the centre of the capital. If the variables chosen had been birthplace (subdivided by, say, Irish, Asian and Afro-caribbean categorizations), then a great many places would have shown either high proportion of all three or low proportions, because these groups of people tend to cluster together in space, although many more subtle mixes of colours would also be visible.

The patterns in *Figure 7* are quite difficult to discern because there is a 'bittyness' which reflects the reality that the social structure is textured. To make those patterns which form more apparent we can generalize the image by applying 'smoothing' operations of generalization (of the type used in remote sensing). *Figure 8* shows the result of using ten passes of the most basic binomial filter (through which, in each dimension, each point is given a new value equal to half its old value and a quarter of the values of each its two closest neighbours (Tobler, 1989). In the figure the very complex distribution of occupational groups has been smoothed. Generalization very much simplifies the image and is particularly useful when reductions in size are required, but, as with all these decisions, the impression of the distribution changes. Watching the generalization take place as an animation can be very enlightening, the image eventually becoming a single blur of one muddy hue. As we generalize like this we are, in effect, using larger and larger (overlapping) areal units – exploiting the modifiable areal unit problem to our advantage.

In *Figure 8* the sharpest divide is confirmed to be in Inner London, where a white buffer appears between the blue epicentre of the capital and its red core. This is where the spatial generalization procedure failed to merge the colours because social divides at this boundary are too great. Moving outwards the red becomes orange, where more "intermediate" workers are housed and then progressively yellow, green and blue as areas housing 'professional' workers are taken in. Purple is the least common colour, as would be expected, representing places where the lowest and highest categories of occupations mix – it is (in various shades), nevertheless, *still present*.

One final word of warning about detailed colour mapping: the apparent colour of an object is affected by the colours surrounding it. Isolated spots of unexpected colour stand out in otherwise uniform areas far more than they do where there is already great variation. The colour of each object you see is, in fact, the result of a mixture of colours on the page that is unique to every distribution shown (Tufte, 1990). This is not such a grave problem when visualizing geographic space as it is with other subjects as we may, anyway, choose to blur adjacent colours for theoretical reasons. We must just understand that our pictures can be as subjective as the conclusions we draw from them – and the decisions which led to their creation – despite the explicit consistency of their design. Each apparent weakness of a choice in map design can always be argued to be a strength!

MAPPING CHANGE AND FLOW

Mapping the census does not necessarily involve only the census, or only one census. Pictures of change over time, and of the spatial relationships between characteristics found in the census and in other surveys often produce more unexpected results. Greater care is also needed when disparate datasets are related graphically. Each census is taken using a unique geographical base and so, at the finest level of detail, censuses cannot be related to one another spatially. Advantage as well as necessity for the use of less fine geographical units can also be argued for; as the sensitivity of the data increases once changes over time are calculated (slight errors in two large numbers result in larger errors once their difference is calculated – Cole (1993) discusses other sources of error).

The enumeration districts (and output areas) of the 1991 census data can be amalgamated to fit, as well as possible, into 1981 wards (and Scottish part-postcode sectors – Atkins and Dorling (1993)). Maps of change over time can then be produced using the 'frozen ward' as a common spatial unit. *Figure 9* shows the change and the new static position in the relative distribution of those people working in 'professional' occupations using a simple five category classification of wards. As the statistics become more complex it is often worth making the map design simpler. The changes are subtle. The static picture of 'professionals' in 1991 at ward level is very similar to the distribution of 'blues' in the 1981 enumeration district cartogram of *Figure 8*. Marked geographical patterns are visible with the most significant increase of professionals being in Inner London, and in other smaller areas of apparent gentrification elsewhere in Britain. *Figure 6* showed that Inner London is also the area where the data is least reliable so some caution in interpretation is advisable.

Although the figures for individual wards might be unreliable, where a cluster of wards have very similar values it is reasonable to assume that, in these places, a significant change has occurred. Central London has obviously experienced substantial gentrification over the decade, whereas other large British cities have not had quite the same experience. *Figure 10* shows the same distribution as *Figure 9* using a cartogram and a traditional

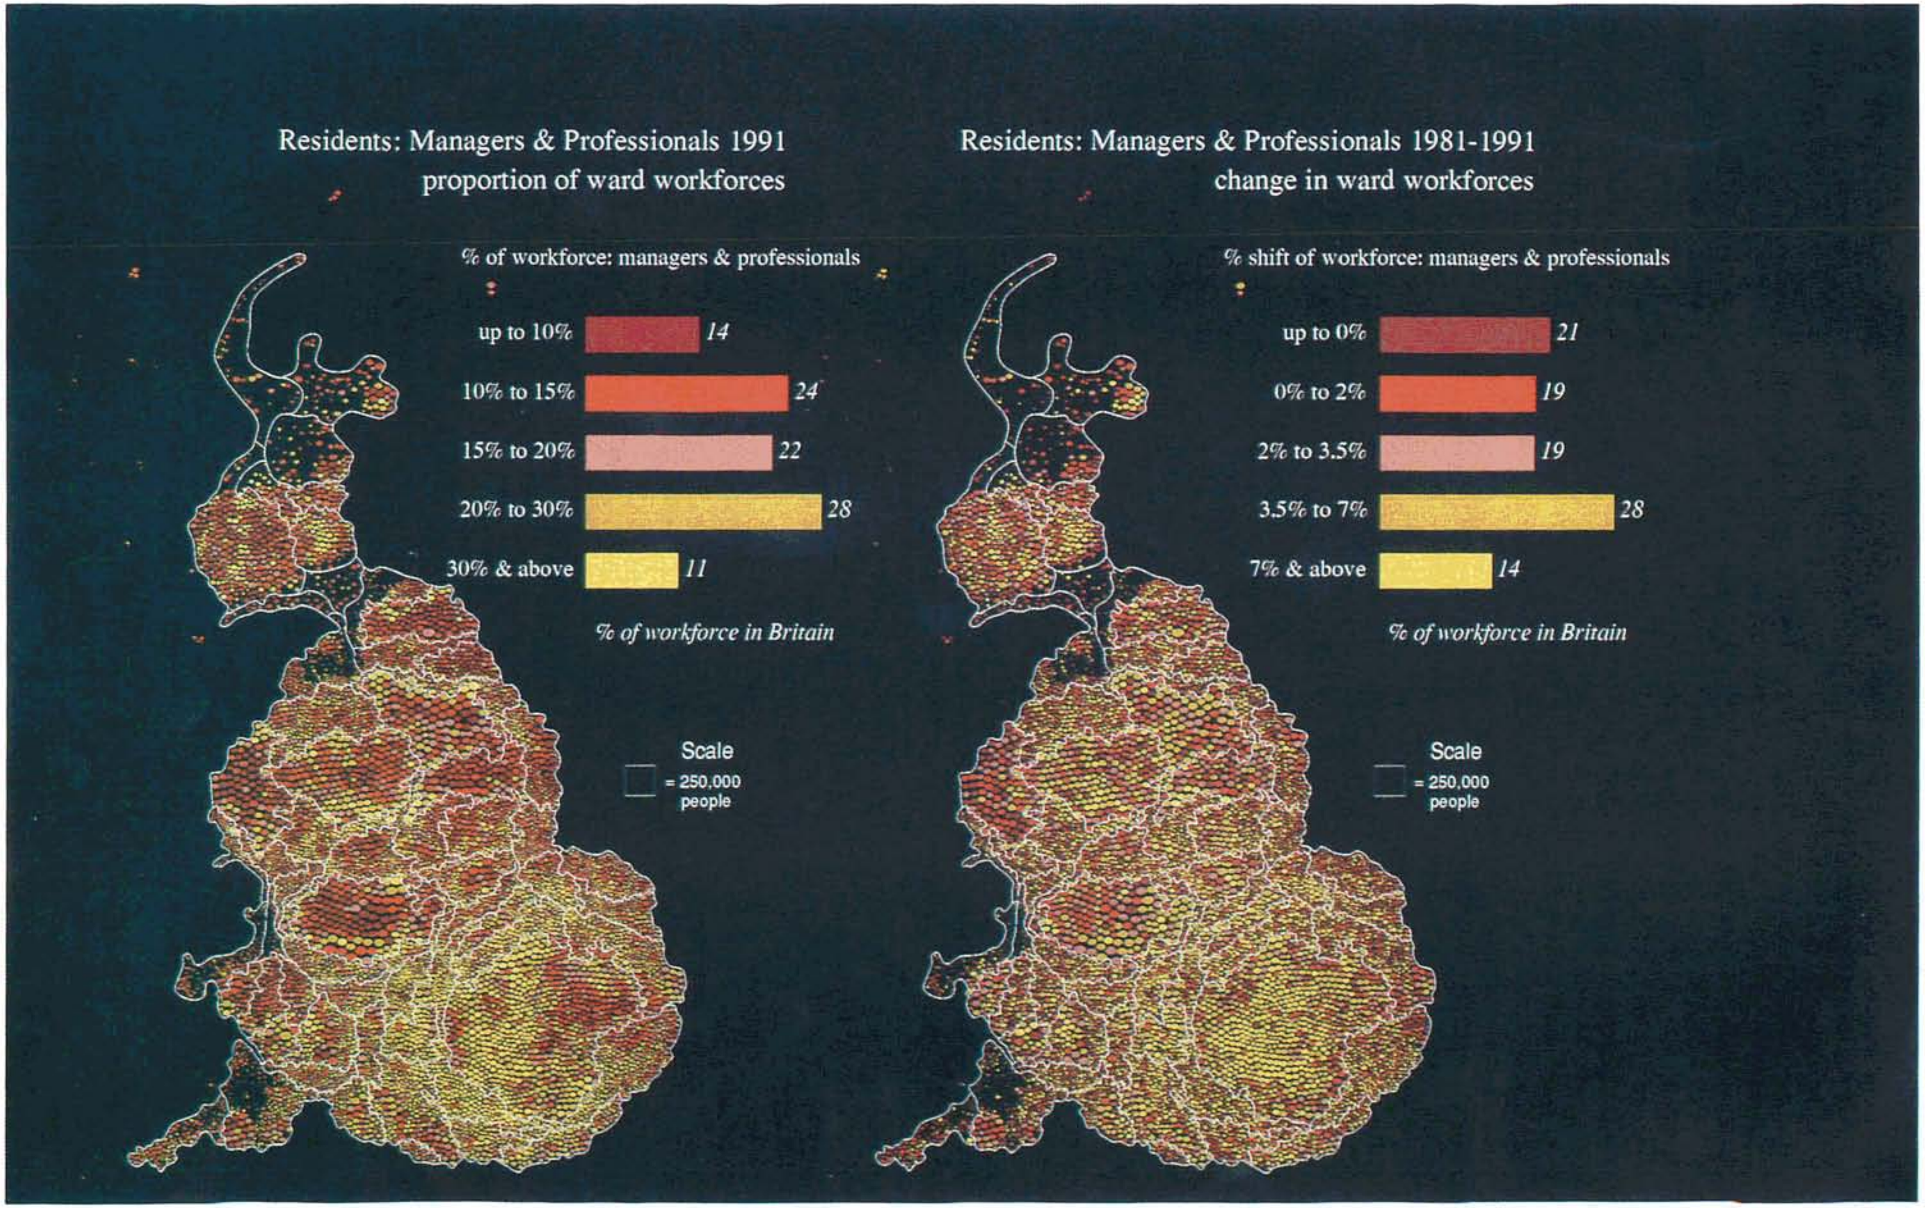

Figure 9. Residents: managers and professionals, proportion and change at ward level.

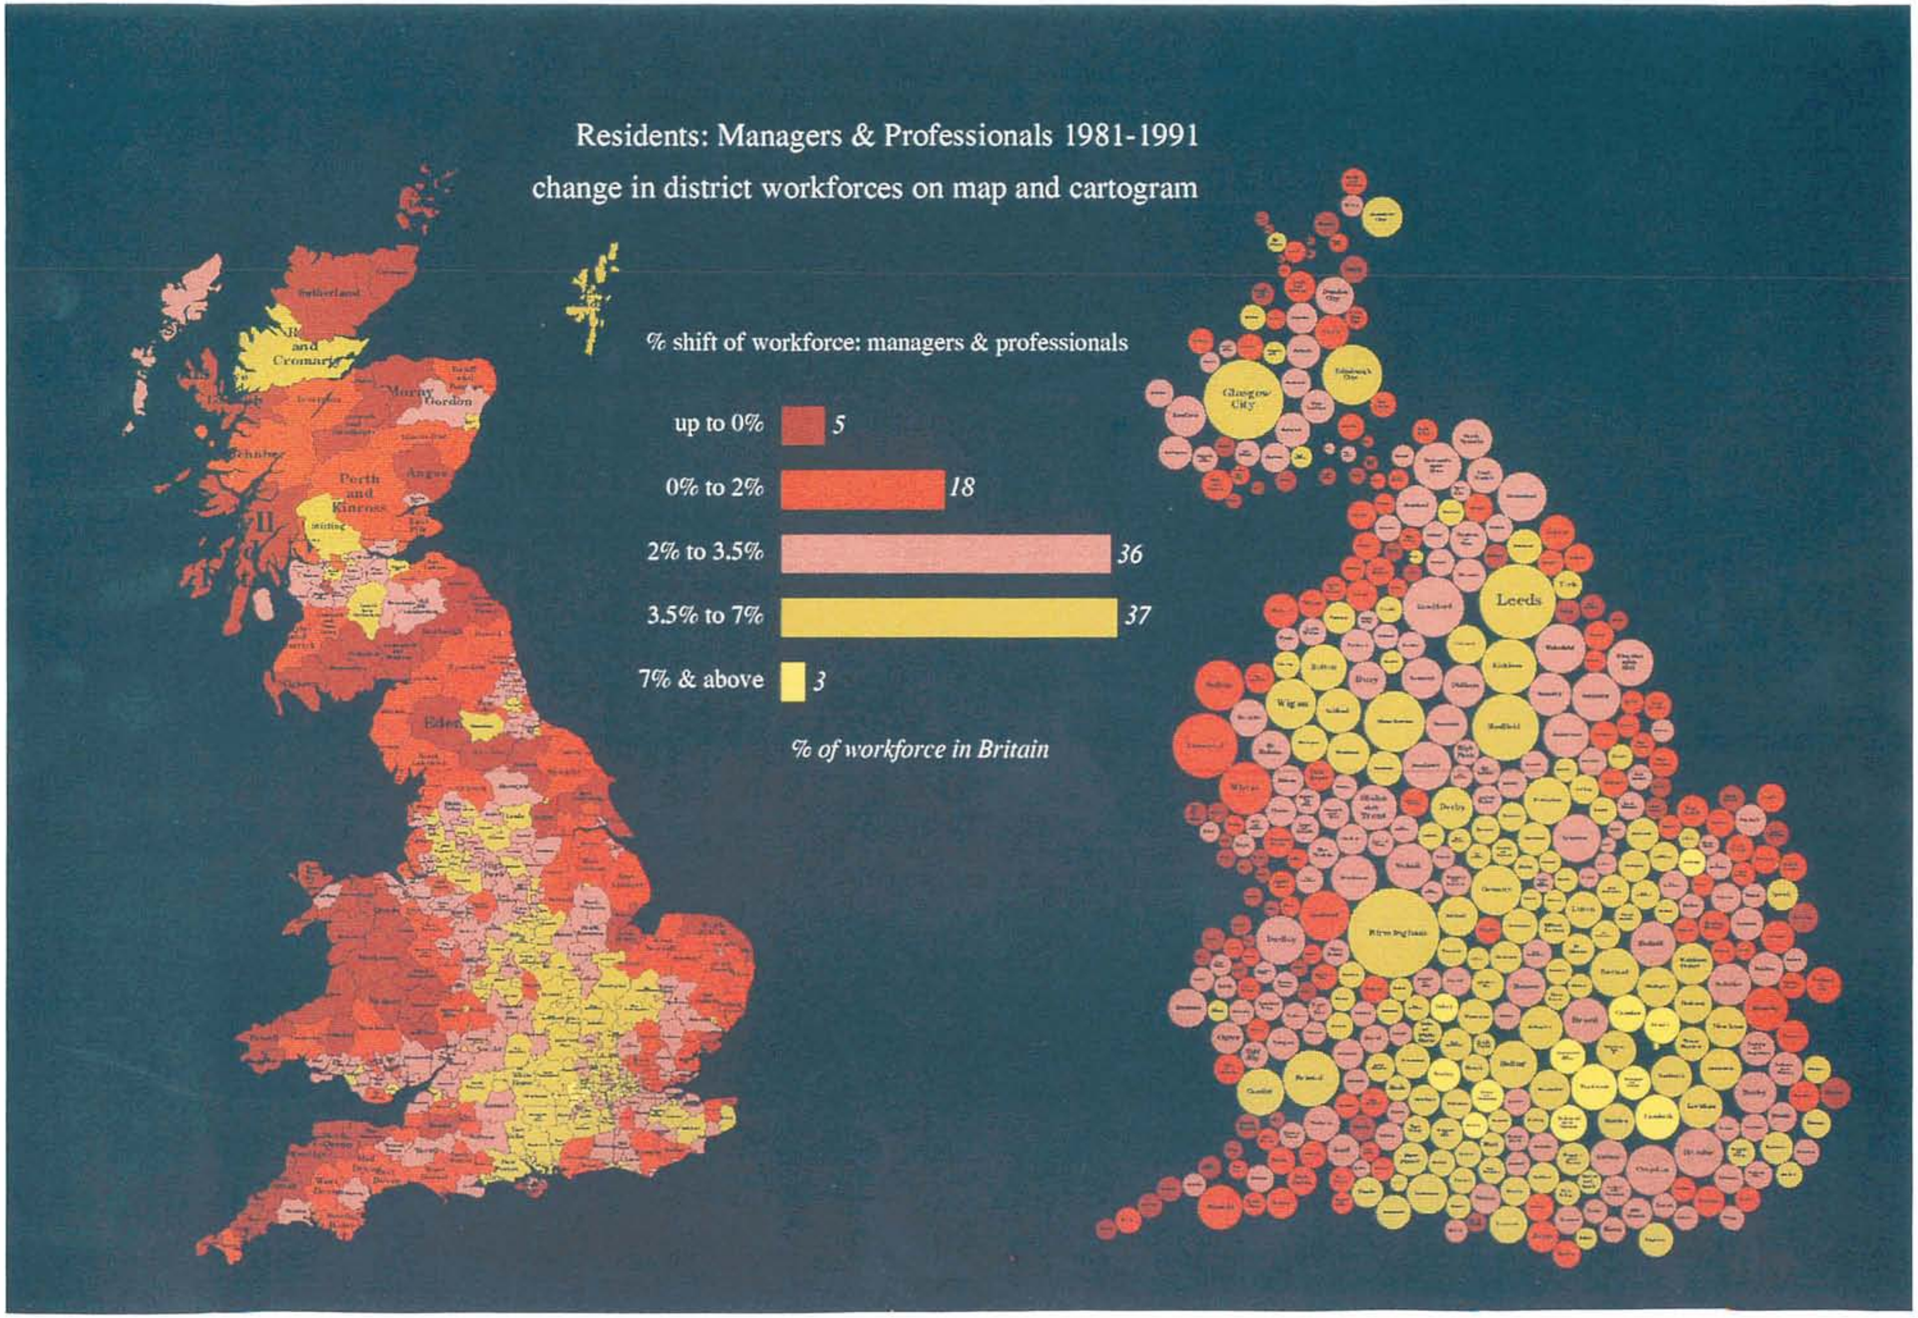

Figure 10. Residents: managers and professionals change on district map and cartogram.

land area map of Local Authority Districts. On these maps place names can be added, as has been done here, although few are likely to be legible when the figure is reduced for printing (the key advantage of desk-top mapping over paper publication is interactive panning and zooming). Here a much simpler spatial story is told, of a central core and declining hinterland to Britain: an easy to absorb tale, but how representative is it? This illustrates how important choices of design are in census mapping – in all senses they totally alter the picture.

Flow mapping is a classic example of the importance of design, and the commuting and migration flow matrices of the census are another set of related datasets to be mapped. These matrices contain information about those people who have crossed any of the boundaries separating over nine thousand areas in England and Wales in the week or year prior to the census date. *Figure 11* shows all the significant migration flows drawn on a traditional map with the lines of flow coloured by the dominant occupational status of the ward of origin. The bright red centres of some cities show that the distances travelled by many people in low status areas are often short. When reprojected on a cartogram the boundaries of Local Authority Districts can be discerned as barriers to people moving between council houses (although more intensive research would be needed to confirm this).

Long distance migration is, in general, the province of those leaving 'professional' or 'professional-intermediate' areas (blue or green/blue lines). Distinct concentrations of flows can be seen from London to the South Coast, and even to the wards of the Isles of Scilly. These may well be due to retirement and a map of flows for every age band could be drawn to discern this. Flow mapping, however, is still very much in its infancy. An area of considerable research challenge is the mapping of spatial change in flow matrices over time. Mapping in both time and space is an area where new technology promises more imaginative methods of visualization.

"The development of census cartography should be based first and foremost on the use of models of the dynamics . . . Look upon the population and its various activities as part of a vertically-rising stream in space-time . . ." Szegö (1987, p. 200).

RELATING OTHER DATASETS

Politics and Health are two areas of concern to social scientists which are closely connected with census mapping. *Figure 12* shows a cartogram of all the wards in Britain, used to visualize the results of the British local elections of 1987, 1988 and 1990 combined (all three years have to be included to get a complete coverage of the country). Each ward has been coloured according to 36 categories of election result, ranging from a Conservative marginal where Liberal came second and Labour third (cyan) to a ward where only a Labour candidate stood (blood red). The standard electoral triangle is used as a key in which the colours of all the possible categories are displayed to indicate the proportions of the vote they represent (Upton, 1991). To explain adequately just this one diagram would take several pages of text. That is not the purpose here: what is of interest for this paper is to show that this detail is possible and to highlight its implications.

The lace-like patterns of local voting tally very closely with those of occupations which were shown in *Figure 8* (as has been known for a long time, but can now be demonstrated visually). Features such as the concentrations of Liberal seats on the edge of Inner London, and the Liberal's general high propensity to appear where there is a greater degree of social mixing, might have been missed in a conventional analysis. The fact that everywhere there are small clusters of bright blue ('safe' Conservative seats), surrounded by purples (Conservative/Labour marginals) encroaching on reds with almost infinite repetition might well be missed without visual analysis – as too can the stark message of left-wing support from most of the electorate in recent local elections.

Moving from politics to health, *Figure 13* shows the spatial distribution of mortality by two of the most significant causes – cardiovascular diseases and cancers. In mapping mortality data accurate census figures are needed to calculate the standardized mortality rates (so that the influence of concentrations of particular age groups and sexes does not dominate the patterns). Here data from both the 1981 and 1991 census have been used along with mortality records from nine of the intervening years. The North/South divide in Heart Diseases and the Inner/Outer City pattern to the spatial distribution of Cancers are both clearly evident. As are, most importantly, all the fallibilities of making such a simple statement about such complex spatial distributions.

CONCLUSION

This paper has only been able to take the briefest look at the new opportunities for map design which the modern census data and computer versatility provides. It has concentrated on high resolution colour mapping because the census is one of the few datasets to provide information sufficiently detailed for such designs. At lower resolutions (when mapping districts and counties for instance) many more graphical opportunities are opened up along with the space in which to implements them (Dorling, 1993).

The census was made for mapping. It contains more spatially referenced information than any other social survey and allows yet more (geographical) information to be meaningfully mapped by providing a base against which this can be done. If we are to do justice to the representation of the people we are claiming to show, then great care is needed in the design of these images. The census contains some seven thousand statistics for each of more than one hundred thousand areas. Combined with other censuses and related datasets access we quickly gain access to billions of numbers, which describe (admittedly, only in the briefest of terms) characteristics of the lives of millions of people. There will never be enough researchers with enough time to look at more than a fraction of this information, but if we can look at it efficiently, fairly and with a better idea of what it is we are looking for and how best to look, we might just learn something new about the society we constitute.

REFERENCES

Arnheim, R., 1970. *Visual Thinking*, Faber and Faber Ltd., London.

Atkins, D. and Dorling, D., 1993. Connecting the 1981 and 1991 Censuses, paper presented at the *1991 Census Conference, Newcastle*, 13–15th September.

Beniger, J. R. and Robyn, D. L., 1978. Quantitative graphics in statistics: a brief history. *The American Statistician*, (32) 1, 1–11.

Bertin, J., 1981. *Graphics and graphic information processing*. Walter de Gruyter, Berlin.

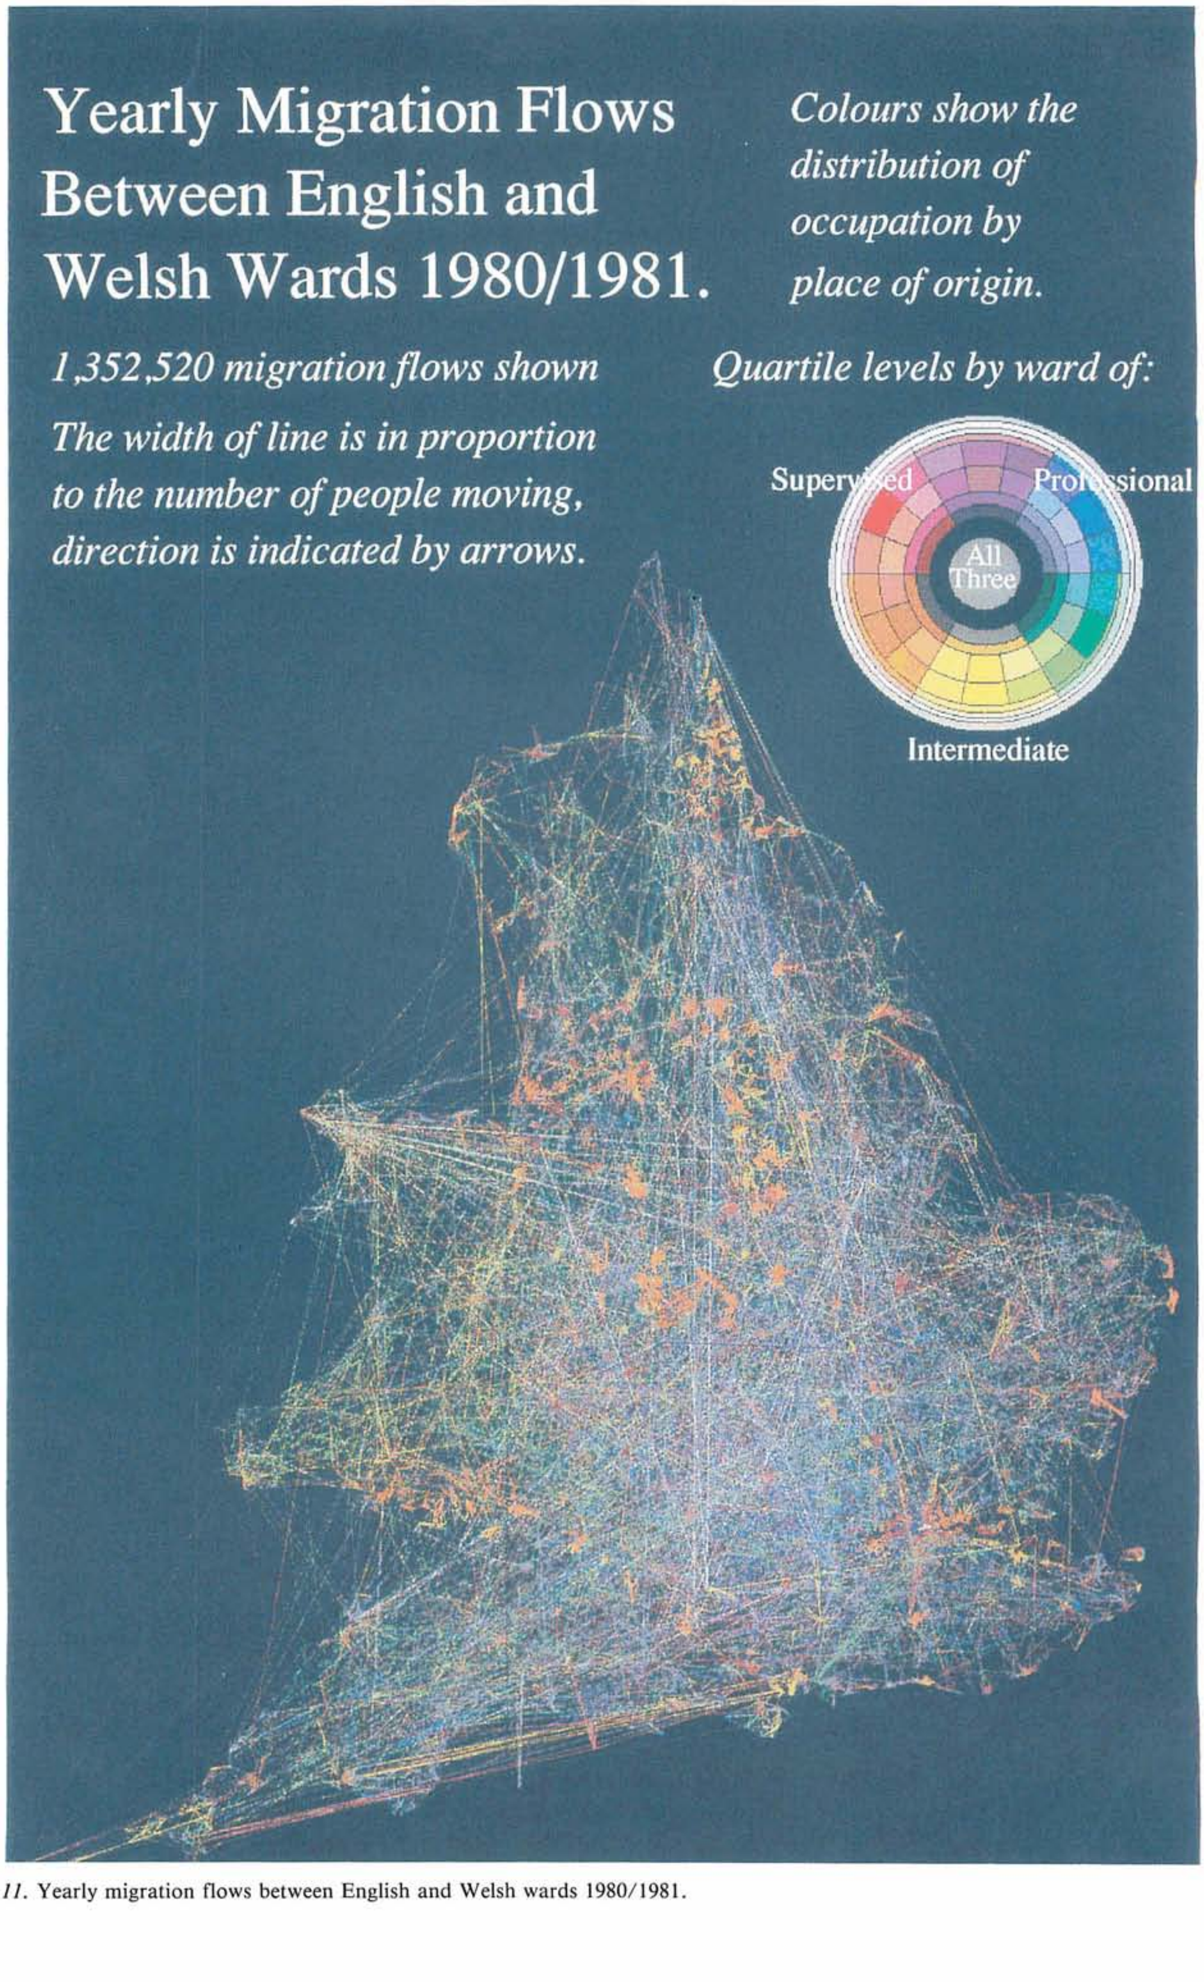

Figure 11. Yearly migration flows between English and Welsh wards 1980/1981.

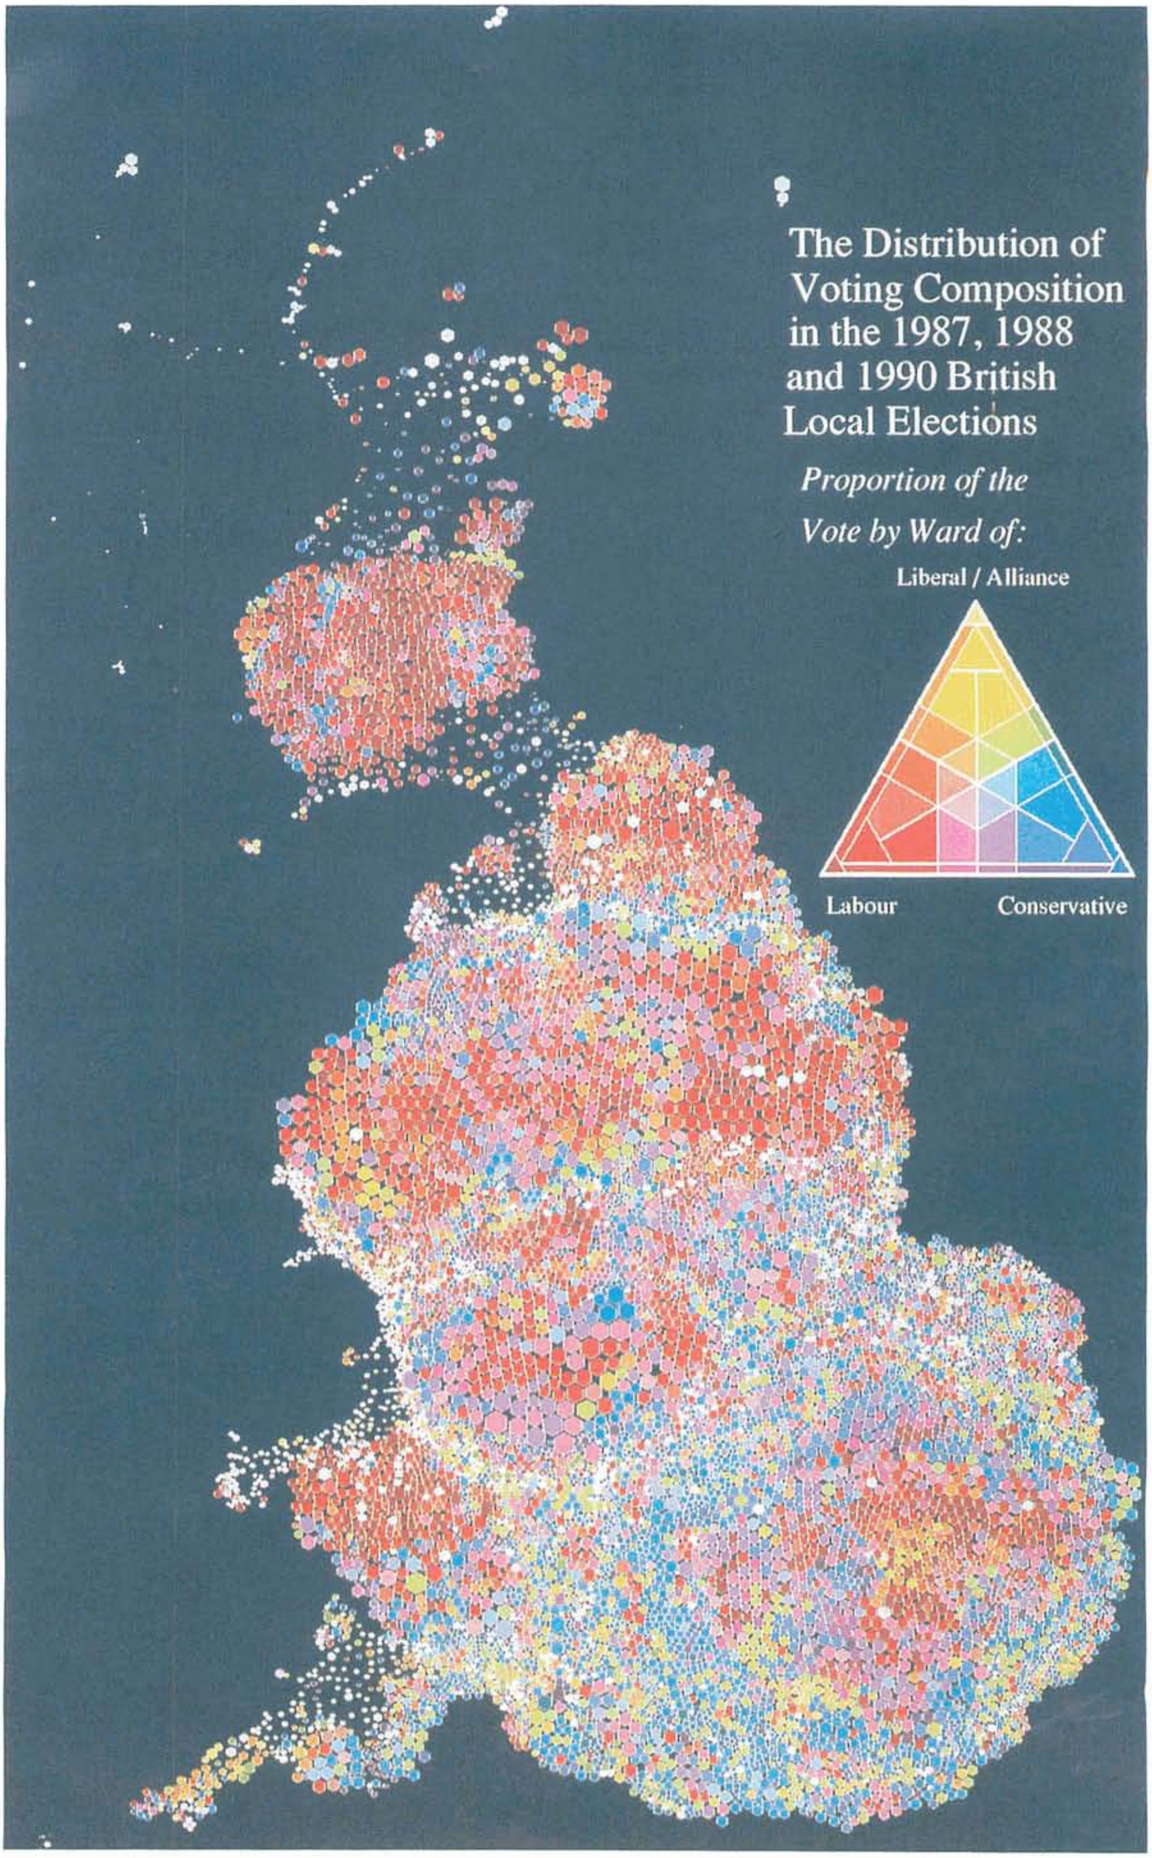

Figure 12. The distribution of voting composition in the 1987, 1988 and 1990 elections.

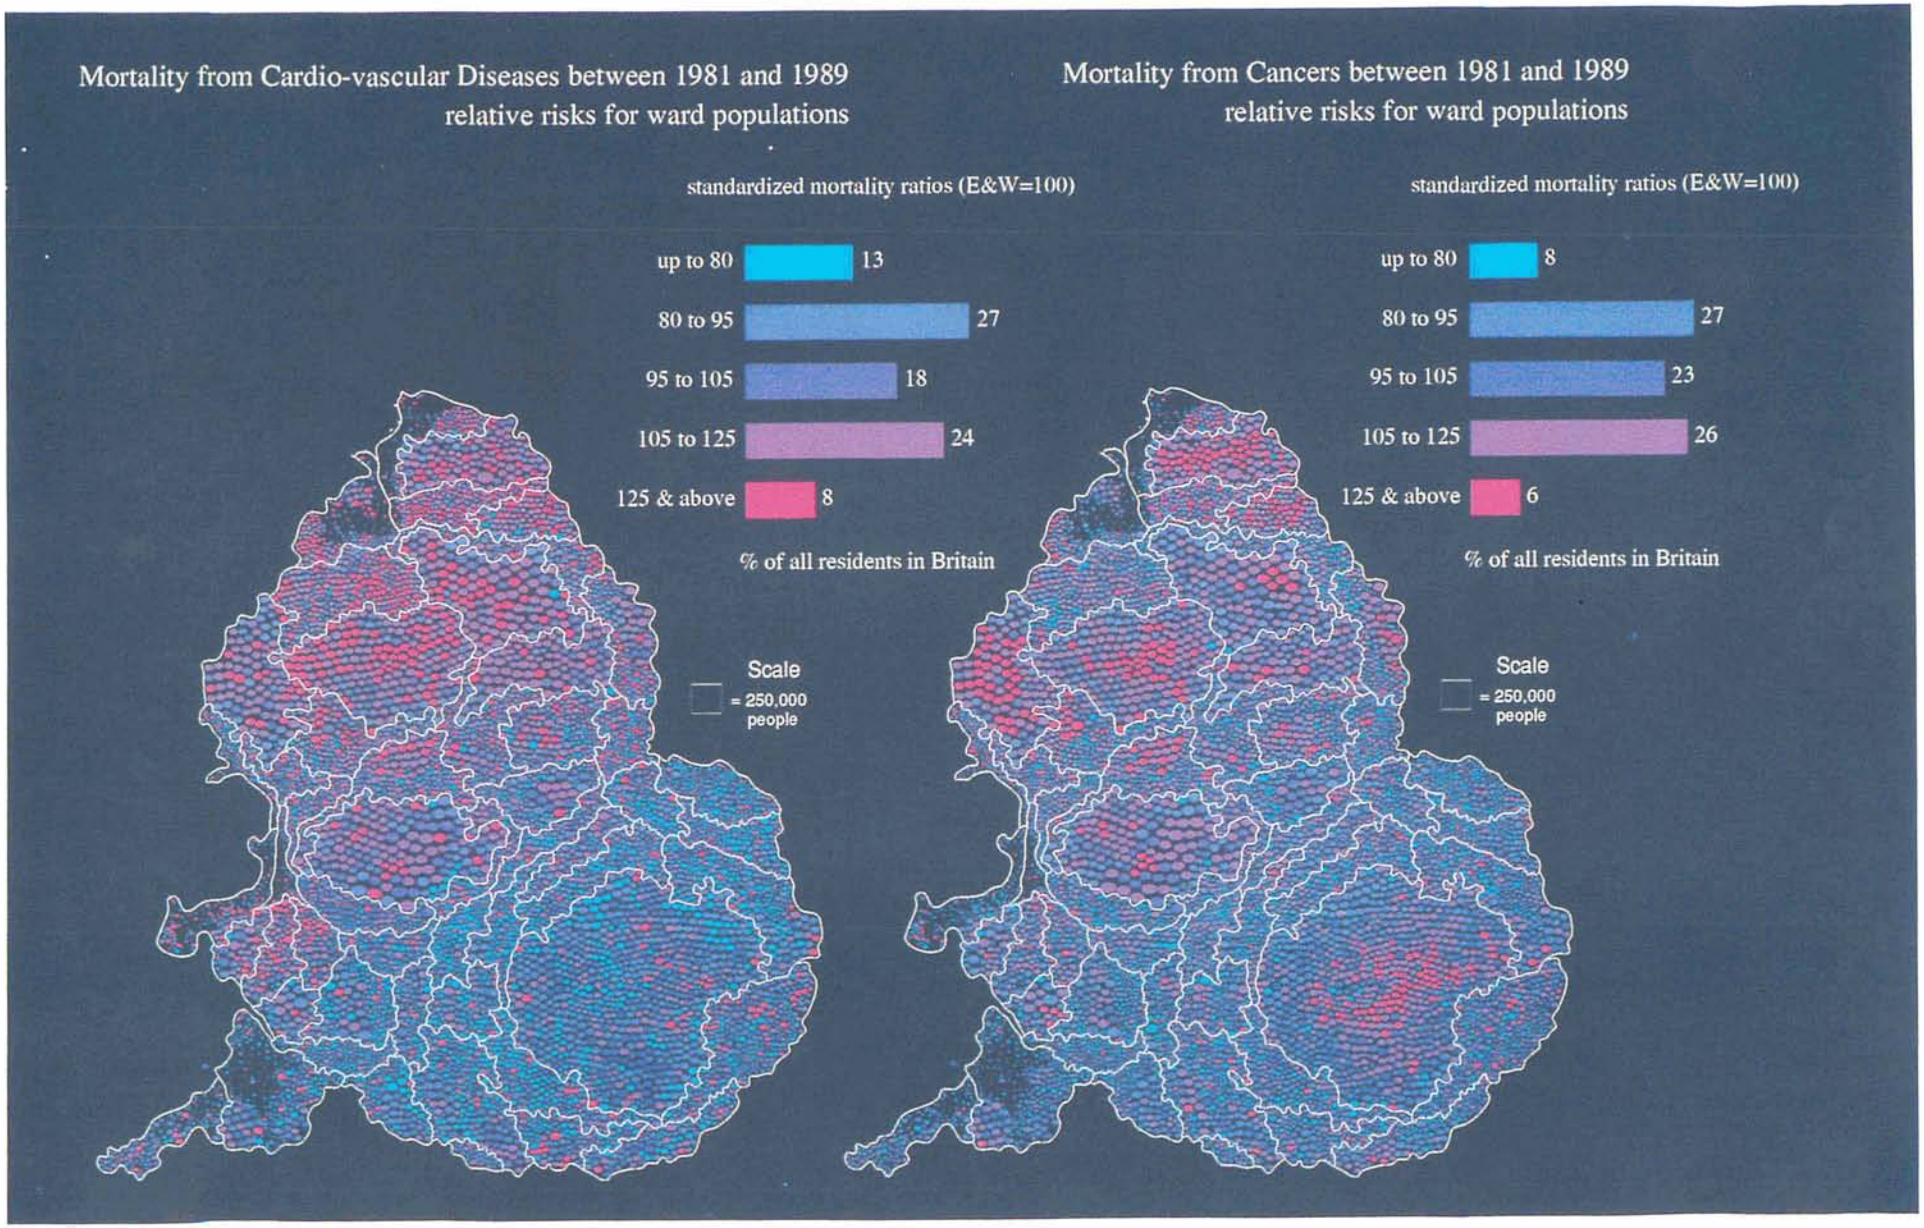

Figure 13. Mortality from cardiovascular diseases and cancers at ward level.

Brunsdon, C., 1991. Estimating probability surfaces in GIS: an adaptive technique. *EGIS '91 Proceedings*, pp. 155-164, Brussels, Belgium.
Census Research Unit, 1980. *People in Britain - a census atlas*. H.M.S.O., London.
Cole, K., 1993. *Sampling error and the 10% small area statistics*. Manchester Computer Centre Newsletter No. 22: University of Manchester.
Dorling, D., 1991. *The visualization of spatial social structure*. PhD thesis: University of Newcastle upon Tyne.
Dorling, D., 1992. Stretching space and splicing time: from cartographic animation to interactive visualization. *Cartography and Geographical Information Systems*. (19) 4, 215-227, 267-270.
Dorling, D., 1993. Visualizing the 1991 census, in S. Openshaw (Ed.). *A Census User's Handbook*. London: Routledge (forthcoming).
Fotheringham, A. S., 1989. Scale-independent spatial analysis. Chapter 19, M. Goodchild and S. Gopal (Eds.). *The Accuracy of Spatial Databases*. Taylor & Francis, New York.
Harley, J. B., 1989. Deconstructing the map. *Cartographica*, (26) 2, 1-20.
Hollingsworth, T. H., 1964. The Political colour of Britain by numbers of voters. *The Times*, October 19th, p. 18.
Howe, G. M. (Ed.), 1970. *National atlas of disease mortality in the United Kingdom*, Thomas Nelson and Sons, London.
Hunt, A. J., 1968. Problems of population mapping: an introduction. *Institute of British Geographers, Transactions*, (43).
Langford, M., Unwin, D. J. and Maguire, D. J., 1990. Generating improved population density maps in an integrated GIS. *Proceedings of the first European Conference on Geographical Information Systems*. Utecht: EGIS Foundation.
Martin, D., 1991. *Geographic information systems and their socioeconomic applications*. Routledge: New York.
Olson, J. M., 1976. Noncontinuous area cartograms. *The Professional Geographer*, (28) 4, 371-80.
Openshaw, S., 1982. *The modifiable areal unit problem*. Concepts and Techniques in Modern Geography No. 38. Gen books: Norwich, England.
Petermann, A., 1852. *England and Wales: distribution of the population*. Census of Great Britain, 1851, Population Tables I. Her Majesty's Stationery Office, London.
Robinson, A. H., 1989. Cartography as an art, in D. Rhind and D. R. F. Taylor (Eds.). *Cartography Past, Present and Future*, 91-102.
Sibert, J. L., 1980. Continuous-colour choropleth maps. *Geo-Processing*, (1), 207-217.
Szegö, J., 1987. *Human Cartography: mapping the world of man*. Swedish Council for Building Research, Stockholm.
Tobler, W. R., 1973. A continuous transformation useful for districting. *Annals of the New York Academy of Sciences*, (219), 215-220.
Tobler, W. R., 1989. An update to numerical map generalization, in R. B. McMaster (Ed.). Numerical Generalization in Cartography. *Cartographics*. University of Toronto Press.
Tufte, E. R., 1990. *Envisioning Information*. Graphics Press; Cheshire, Connecticut.
Upton, G. J. G., 1991. Displaying election results. *Political Geography Quarterly*, (10), 3, 200-220.
US Bureau of the Census, 1970. *Distribution of older Americans in 1970 related to year of maximum county population*. United States Maps, GE-70, No. 2. US Department of Commerce, Washington DC.
Visvalingham, M., 1981. The signed chi-score measure for the classification and mapping of polychotomous data. *The Cartographic Journal*, (18), 1 32-43.

ACKNOWLEDGEMENTS

Thanks are due to Tony Champion, James Cornford, David Dorling and Nina Laurie who read and commented on earlier drafts of this paper, and for funding from the British Academy. Joseph Rowntree Foundation and The University of Newcastle upon Tyne which enabled the research this paper is based upon to take place.