RESOURCE ALLOCATION USING MEASURES OF RELATIVE SOCIAL NEEDS IN GEOGRAPHICAL AREAS:

The relevance of the signed chi-squared, the percentage, and the raw count.

Steve Simpson, City of Bradford Metropolitan Council
and Census Microdata Unit, University of Manchester

MANCHESTER

CENSUS MICRODATA UNIT

1994

ISBN 1 899005 05 6

Printed in Great Britain by:

Census Microdata Unit
Faculty of Economic and Social Studies
The University of Manchester
Oxford Road, Manchester, M13 9PL

RESOURCE ALLOCATION USING MEASURES OF RELATIVE SOCIAL NEEDS IN GEOGRAPHICAL AREAS:

The relevance of the signed chi-squared, the percentage, and the raw count.

CONTENTS

Resource allocation using measures of relative social need in geographical areas: the relevance of the signed chi-squared, the percentage, and the raw count.

Stephen Simpson, City of Bradford Metropolitan Council and Census Microdata Unit, University of Manchester

Summary

The construction of social indices to compare local areas is an important influence on the allocation of public resources. There has been debate in past years on the choice of variables for such indices and the means used to combine percentages measured for more than one variable. Recently the percentage and its z-score have given way to the signed chi-squared statistic as an apparently more appropriate means of identifying concentrations of need. The relative merits of the signed chi-squared, the percentage, and the raw count are discussed here in relation to two separate objectives in resource allocation: identification of concentrations of need (geographic ring-fencing) and calculation of local budgets (resource distribution). The signed chi-squared is shown to overcome disadvantages of the percentage when identifying high concentrations of need, but to entail its own disadvantages. It is arbitrary in its construction, its ranking of areas is dependent on the choice of a reference area, and it is misleading for areas of low concentration. These conclusions apply both to comparison of large local areas such as local authority districts, and to comparison of small areas within a locality. Alternatives are discussed.

1 Introduction

1.1 Two stages in geographic resource allocation

When responsibility for public services is devolved to geographical area budgets, two important decisions have to be made: to which areas should resources be directed, and how much does each chosen area merit?

Some budgets are restricted geographically. For example, the public grants given to those areas with Assisted Area status, or with European Community status I, II, III, or IV; similarly, a local authority may direct a variety of community regeneration programmes to selected communities of concentrated need. The process of identifying the areas which merit targeted attention is termed here 'geographic ring-fencing'. Geographic ring-fencing is usually a response to perceived social inequality, however defined, between these areas and others, and is an attempt to direct resources to lessen such inequality or to dilute the concentrated need. Geographic ring-fencing is preliminary to the task of dividing a total budget among the receiving areas, here termed 'distribution'.

In some cases it is clear that all areas will receive some share of the total budget. For example in the calculation of Standard Spending Assessments used to decide the core national grant for local councils in England, or in the devolvement of social services budgets to decentralised cost centres within a local authority, or in the calculation of individual school budgets. In all these cases, there is no geographic ring-fencing to reduce the number of areas or institutions that receive a share of the total budget.

These decisions about resources, of 'where' and 'how much', are often made on an historical basis, by adjustments to existing allocations. Within a commercial market for services, the placing of new outlets and the size of local budgets would react to demand for the service. However, public expenditure is often directed to those whose income is insufficient to create a commercial market for services. In addition, the past takeup of services is often not a good guide to need, and is not an available guide to new services. Then, as in the examples above, an effort may be made to identify objectively areas of relatively concentrated need in order to ring-fence budgets to those areas, and to measure objectively each area's relative aggregate need in order to distribute budgets fairly between them. Funding on the basis of estimated relative need is often termed 'formula funding'.

There has been considerable discussion of the most appropriate indicators of need for the service in question, and of the ways in which indicators should be standardised, transformed and combined to provide an overall index of need for the service. Useful summaries of that discussion are provided by Colin Thunhurst (1985), Martyn Senior (1991) and Mike Derbyshire (1993) and in a collection of conference papers from local authorities and the Manchester Census Group (Simpson, 1993). Some reference will be made to these issues, but this paper focuses on a less aired but perhaps more fundamental debate associated with the estimation of need, the unit of measurement.

1.2 Three units for estimating relative need

It has usually been assumed that once a variable (such as unemployment) has been chosen to represent need or an element of need, then the percentage (the unemployment rate) is the correct 'unit' in which to measure relative need. That is, areas with higher percentages correctly identify areas with greatest need. Recently the supremacy of the percentage has been questioned.

On the one hand it has been argued by Keith Cole (1993) and others that an alternative measure is more appropriate for identifying concentrations of need, the signed chi-squared statistic which weights the area percentage by the population in the area. A quantitative study of urban poverty in English local authority districts commissioned by the Department of Environment used the signed chi-squared as its basic unit of measurement (Bradford et al, 1994); the study, from the Manchester Centre for Urban Policy Studies is used for the geographic ring-fencing of resources involved in the Single Regeneration Fund which replaces the Urban Programme and other funds (Government Offices for the Regions 1994: 10).

On the other hand the simple count of an indicator is often used in practice instead of the area percentage, for distribution of funds to devolved budgets (Simpson 1993: 6-7; Browne 1992; Jenkins 1994).

The aim in this paper is to discuss each of these three basic units in turn - counts, percentages and signed chi-squared statistic - in relation to the two major budgetary purposes of social needs indices - geographic ring-fencing to areas of concentrated need and distribution of resources according to aggregate need in each area.

2 Counts

2.1 Distribution of a budget

Here the number of people affected by a relevant condition in each local area is taken to reflect the aggregate need for appropriate resources in that area. The number of children in a local area whose households have no earner might be taken to reflect the relative need for priority nursery places. The available total budget (after possible deductions for some historical committed resources and for a practical base budget in each area) is then divided between local areas according to the number of such children found in each area.

The principle extends readily to more than one indicator, simply requiring that each indicator I_i is given a weight W_i, which is in effect the proportion of the total available budget B which will be allocated according to this indicator. Then area x will receive a budget B_x calculated according to the raw counts I_{ix} in that area:

$$B_x = B \times \left(\sum_i W_i \frac{I_{ix}}{\sum_x I_{ix}} \right) \qquad (1)$$

Table 1 gives an hypothetical example using the indicators children in low-earning families, in crowded households, and in lone parent families.

Table 1: Distribution of resources: worked example

Indicator	Weight W_i	Raw counts: I_{ix} Area1	Area2	Area3	ΣI_{ix} Total	Counts after scaling: $I_{ix}/\Sigma I_{ix}$ Area1	Area2	Area3	Total
Low earners I_1	0.5	1,317	912	304	2,533	.52	.36	.12	1
Overcrowded I_2	0.3	354	303	355	1,012	.35	.30	.35	1
Lone parents I_3	0.2	254	203	178	635	.40	.32	.28	1
	1.0								
		Composite index, with weights W_i:				.445	.334	.221	1

(Area 1 has 52% of the children in families with no earner (100*1,317/(1,317+912+304)), 35% of those in overcrowded housing, and 40% of lone parents. Its composite index is 0.5*.52 + 0.3*.35 + 0.2*.40 = .445; this is the proportion of the total budget which is suggested it merits.)

The raw count for an appropriate indicator is therefore quite usable for the distribution of a budget between chosen areas. Formula (1) is commonly used by local authorities. Martin Browne (1992) describes the indicators and weights used by five local authorities' social services departments to distribute their home care budgets to area offices.

The choice of indicators and their weights relies on the indicators' relationship to need for the service in question. Sometimes this relationship can be quantified by special studies as advised by Martin Browne (and achieved by John Jenkins (1994) in the case of child care in Cambridgeshire). Hugh Davies, Heather Joshi and Lynda Clarke (1995) have attempted to find the weighted combination of census indicators that best predicts low income. In practice however, the indicators and weights are often provided by no more than a pragmatic consideration of social theory, and the judgement of the producers of the index and the managers of the budget.

In the examples quoted by Martin Browne, historical patterns of demand are frequently used as one variable in the index, tempering indicated need by past practise.

2.2 Geographic ring-fencing

Raw counts are not suitable for identifying areas with high concentrations of need. Think of two areas each with the same count of those in need, but with area A having five times the population of area B. If the focus is to identify areas of concentrated need in order to address inequality of access to services, then area B has the better case since a much higher proportion of its residents are in need.

Thus we are led to consider the percentage of each area's residents who are in need. This uses as denominator the area's own population, rather than the sum of those in need across all areas as in this section on counts.

3 Percentages

3.1 Distribution of a budget

The percentage of a relevant population that is indicated to be in need is the most common unit for comparing the degree of need in different areas; for example the percentage of all residents who are under five, or the unemployed expressed as a percentage of all economically active residents (those unemployed *and* those working). When various indicators are combined into a single index of need, each variable is usually first standardised as a z-score so that all have equal mean and variance. The discussion in this paper of the percentage applies equally to its standardised form and to other ratio statistics such as Standardised Mortality Ratios.

The percentage is a measure of the average need of individuals in an area. By its construction the percentage ignores the relative size of each area.

Part (a) of table 2 shows the count of the numbers of households without a car in two districts of England, and that count expressed as a percentage of all households. The level of car ownership is a commonly used indicator of average income in an area. While Leeds has greater aggregate need as displayed by the high count of households without a car, this information is lost when the percentages are compared.

When an index of need based on percentages is used to distribute resources, the index therefore indicates a per capita amount that is then multiplied by the area population again in order to calculate the aggregate need deserving of resources in the area. So for example the Jarman score for GP workload is multiplied by the patient list size to calculate the extra resources for GPs in priority areas (Senior, 1991). The Standard Spending Assessment social needs indices are multiplied by the size of the population to calculate a part of the amount of expenditure that central government deems each local authority needs to spend (Chipping, 1994; Senior, 1994). In distributing a budget for schools, the local education authority calculates a per capita amount that may vary according to the special needs of each school and its area, and multiplies it by the number of pupils in the school.

Table 2: Count and percentage, an illustration.

	Count of households without a car	All households	Count as a percentage of all households
(a) Two districts:			
Great Grimsby	15,522	35,419	43.8%
Leeds	116,134	280,845	41.4%
(b) Constituencies within Leeds district:			
Elmet	9,720	34,881	27.9%
Leeds Central	21,015	33,331	63.0%
Leeds East	17,122	32,934	52.0%
Leeds North East	11,855	34,381	34.5%
Leeds North West	11,130	32,707	34.0%
Leeds West	17,617	35,731	49.3%
Morley and Leeds South	14,147	32,818	43.1%
Pudsey	10,827	36,093	30.0%

For the distribution of a budget, this use of an indicator or index of need based on percentages, applied as a per capita amount to the relevant population, appears reasonable. One minor drawback can be noted. The indicators should also refer to the relevant population, but this is often not achieved in practice. For example, the unemployment rate measured in the Jarman and other indices (a percentage of the economically active) is then applied to the whole population, while in fact the economically active is a variable part of the population. The more direct count-based indicators in section 2 avoid this misleading bias when estimating aggregate need.

3.2 Geographic ring-fencing

When used to identify areas of concentrated need in order to geographically ring-fence resources, the percentage has a greater drawback due to the spatial heterogeneity of most social variables. For a large area, the average level of need reflected in the percentage will usually hide higher levels of need within substantial sub-areas. Thus part (b) of table 2 shows that Leeds has areas within it as large as Great Grimsby but with a higher percentage of households without a car. A search for concentrated need in England should identify Leeds as having similar or greater concentrations than Great Grimsby, but could not do so on the basis of district percentages alone.

This means that smaller areas are at an advantage if the percentage is used to identify concentrations of need, as is the case with two of the major social needs indices of the 1980s: the Jarman index already referred to, and the identification of Assisted Areas in government urban policy of the 1980s (DoE 1983). Percentages hide the heterogeneity of larger areas even when it is a collection of relatively small areas that are being considered. Both Keith Cole (1993) and Sally Holtermann (1975) point out that when an index is calculated for the smallest census output areas in Britain (Enumeration Districts in England and Wales, Output Areas in Scotland), Scotland provides more than its fair share of the high-concentration areas simply because its average size of output area is considerably smaller than in the rest of Britain.

Apart from this major difficulty in interpretation of the percentage as a measure of concentration in areas with different populations, there is a problem of unreliability of data for small areas. The effects of measurement error in any data source (and sampling error in some) makes comparison involving small areas an unreliable guide to real differences. In the case of the census a specific additional source of unreliability is the modification of the Local Base and Small Area Statistics to help protect confidentiality of individuals represented in the tables. The tables for areas within local authority districts are adjusted by the random addition 1, 0 or -1 to non-zero counts, in both the 1981 and the 1991 census. This modification of the data, can accumulate within a table such that published totals differ from their measured value by more than one. Cole (1993) describes the modification in detail and points out that even for a count that can only have been modified by one, percentages for a small enumeration district can be drastically affected. Thus a tabulation showing two out of 22 households lacking exclusive amenities may in truth have from 4.5% to 11.7% of households without amenities, simply as a result of uncertainty as to whether the numerator is in fact 1, 2, or 3. As the data modification is applied regardless of the size of the true cell count, the impact on percentages is less for larger areas.

Thus as a means of identifying concentrations of need, the percentage suffers from two serious problems: the greater heterogeneity within larger areas, and the lesser reliability of data for small areas. One approach to address both these problems is to use the percentage in a way that gives greater weight to areas with large populations. The signed chi-squared is based on this approach.

4 Signed chi-squared

The signed chi-squared statistic has been proposed only for geographic ring-fencing to identify those areas which have concentrations of any particular characteristic, and not for use in calculating the aggregate need to spend in each area. Thus the structure of this section is different from previous sections.

The signed chi-squared is not common within the literature of social indicators. This section addresses its origins (including the justifications of those who have used it, and its calculation), and with the aid of examples it explores some of the properties of the chi-squared statistic that will help clarify its relevance in particular applications.

4.1 Origins

The Census Research Unit at the University of Durham advised the use of signed chi-squared when mapping census social indicators (Rhind 1983). It was suggested specifically to deal with the greater unreliability of percentages in small areas due to the data modification by census offices as has been described above.

"Giving graphic expression to the uncertainty in choropleth maps is rather more difficult and it is best to compute a new statistic to map which takes this into account. One such candidate is the signed chi-squared measure ...This takes into account both absolute and relative (eg percentage) deviations from some stipulated average figure such as the national average. ... To be included in the outermost categories, the area must be of a low population but have an enormous proportionate deviation from the nominated mean value, or alternatively, have a large population but a much smaller proportionate deviation from the mean: thus to be in an outer class signifies that the area is distinctly unusual and needs further, detailed examination." (Rhind 1983: 185)

Rhind also argued that even without modification, small areas were more likely to include extreme values (the heterogeneity argument), and thus the chi-squared statistic was a useful "compromise measure ... [whose] results are usually more readily interpretable than those of either ratio or absolute-number based maps".

The major pre-1991 use of signed chi-squared was indeed in the mapping context and from the Durham Census Research Unit, in the volume of maps from the 1971 census results produced for the census offices (CRU/OPCS/GRO(S) 1980). This volume has a usefully clear description of the derivation and interpretation, in the mapping context, of the signed chi-squared statistic (pp.6-7 and pp.124-5).

Following the 1991 Census, the Centre for Urban Policy Studies at the University of Manchester has used the signed chi-squared statistic in preference to percentages for measuring the census and other variables in its Index of Urban Conditions, commissioned by the Department of the Environment to identify areas requiring priority attention in its urban development programmes. Its authors describe their choice of chi-squared as follows:

"The familiar z-score based on suitably transformed data was used for the equivalent index for 1981. This is, however, inappropriate for enumeration district (ED) data where the denominator is small and variable, and the numerator may be affected by Barnardisation [data modification as described above]. A signed chi-square value is more appropriate. It measures the extent to which the ED value deviates from the national value. It lends more weight to EDs with larger denominators and zero reflects the value for the nation. At the district level it was also employed partly for consistency, but mainly because it reflects absolute as well as relative amounts of deprivation, something that one statistic rarely does." (Bradford et al, 1994)

The chi-squared statistic has its origins within statistical theory. In the social context, it compares the count of people in each of a number of categories (perhaps age-groups or employment statuses) as observed in a random sample from a survey, to the counts expected from a prior hypothesis about the population from which the sample was drawn. The chi-squared statistic summarises the set of comparisons between observed and expected values (O-E) in each category i:

$$\chi^2 = \sum_i \frac{(O_i - E_i)^2}{E_i} \qquad (2)$$

Variables used in social needs indexes have just two categories, such as unemployed/not unemployed, not car-owning/car-owning. The differences O-E in the two categories can then be expressed in terms of the sample percentage p with the characteristic, n the denominator of that percentage in the sample, and the percentage π expected by the prior hypothesis. The chi-squared statistic of (2) readily reduces to the following expression, which is also more suitable for computer calculation:

$$\chi^2 = \frac{(p-\pi)^2 n}{\pi(100-\pi)} \qquad (3)$$

Since $\pi(100-\pi)/n$ is the squared standard deviation of the difference between a sample percentage p and its expected value, (3) shows that the chi-squared statistic standardises a percentage by its own standard deviation. It is thus a way of dealing with the unreliability of percentages for smaller samples. In that context, the statistic is helpful in determining how often sample values of the percentage p may differ by large amounts from the population value π; for example the chi-squared statistic exceeds the value of 3.84 in only 5% of random samples.

The chi-squared statistic is indeed a product of both the sample percentage (p) and the observed number (np/100), each expressed as a deviation from their expected value, as can be seen from writing (4) in a slightly different fashion:

$$\chi^2 = \frac{(p-\pi) \times (np-n\pi)}{\pi(1-\pi)} \qquad (4)$$

Now we return to social indices, where the chi-squared indicator has been borrowed from statistical theory because it weights the percentage by the count of a characteristic. In social indices, the random sample is a geographic area, p the observed percentage in the area, and π the percentage in some relevant comparator, for example the nationally observed percentage.

However, the statistical theory does not apply to social indices, for two particular reasons. In the first place, the nationally observed percentage is not a mean around which areas randomly fluctuate. Areas are not random samples of individuals. Indeed, resource allocation formulae under discussion are only required because the distribution of social need is shown, empirically, to be clustered. In the case of car-ownership for example, the chi-squared statistic exceeds 3.84 in 98% of districts of England, far more than the 5% that would be the case if areas were random collections of individuals.

Secondly, in order to distinguish areas above and below the reference area, a negative sign is given to the chi-squared statistic for areas where p is less than π: thus the 'signed chi-squared statistic', χ^2_s. The sign is a creation of social geographers rather than of statistical theory.

The signed chi-squared is thus an appropriation from a sampling theory that is not particularly relevant to the social indicators used for resource allocation. It is simply a convenient weighting of raw counts and percentages, measured as deviations from a standard population. The weighting does not necessarily reflect the reliability of the calculated percentage for each area, nor the likelihood of finding concentrations of need within a large area. An examination of the properties of the chi-squared through some examples will allow readers to judge whether it is a helpful statistic in particular applications.

4.2 Examples, and properties of the signed chi-squared statistic

Table 3 gives the chi-squared statistic, again for Leeds and Great Grimsby districts, for households without a car and for residents in households that lack exclusive use of inside toilet or a bath/shower, another often-used census social indicator. Although the number of such households has decreased in recent decades, precisely due to targeted improvement programmes, it is still used in many social indices. For each indicator, the percentage for England is taken as the reference percentage π, ie 32.4% of households without a car, and 1.09% of residents in households lacking the designated exclusive amenities.

For households without a car, the signed chi-squared takes a similar value in both districts. It gives weight both to the number of households lacking a car which is higher in Leeds, and the same phenomenon as a percentage of the total number of households which is higher in Great Grimsby.

Leeds and Great Grimsby have less than the English norm of households lacking amenities, and so their chi-squared values are both given a negative sign for this indicator. Leeds' greater population means that its chi-squared statistic is much further below zero. But Leeds has a higher count *and* a higher percentage lacking amenities and so it would be hard to argue that it has less concentrated need than Great Grimsby. This is perverse behaviour in the context of resource allocation (although arguably not so in the context of mapping where large negative values will stand out as differences from the national average). Other implications of negative values of the signed chi-squared statistic are illustrated below.

Table 3: Count, percentage and chi-squared, an illustration

	Leeds	Great Grimsby	England
Number of households (n)	280,845	35,419	18,683,337
Numbers of households without a car (O)	116,134	15,522	6,058,602
As a percentage of all households (p)	41.4%	43.8%	32.4%
Signed chi-squared (χ^2_s)	698.04	686.14	0
Number of residents in households (n)	672,769	89,389	46,337,368
Residents in households lacking amenities	4,090	490	503,194
As a percentage of all residents (p)	0.61%	0.55%	1.09%
Signed chi-squared χ^2_s)	-1,431.0	-240.7	0

Figures 1 and 2 show both the percentage and the signed chi-squared for the same two social indicators, for the 366 districts of England (including London Boroughs), with an indication of the size of each district. In the context of resource allocation, two clear advantageous properties of the statistic are evident from figures 1 and 2:

- *Areas with percentage above or below the reference area have positive or negative signed chi-squared values respectively, areas with the reference area percentage have a chi-squared value of 0.*

This easy relation between the reference area and the signed chi-squared value is useful in its interpretation.

- *For a given percentage above the reference percentage π, larger areas score more highly.*

This is the clearest advantage of the chi-squared statistic for identifying geographical concentrations of need. It gives weight to our knowledge that the percentage for a large area is likely to mask sub-areas with higher concentrations of need. It also gives an allowance for the greater unreliability of census data for the smaller areas such as Census Enumeration Districts that are subject to data modification as discussed above.

A third advantage of the chi-squared statistic over the percentage is illustrated by figure 3 and table 4, which show the cumulative percentage of poor individuals included in the 'poorest' districts when ranked by the percentage and by the chi-squared statistic.

Figure 1: Percentage and signed chi-squared statistics,

Household without access to a car

366 Districts of England

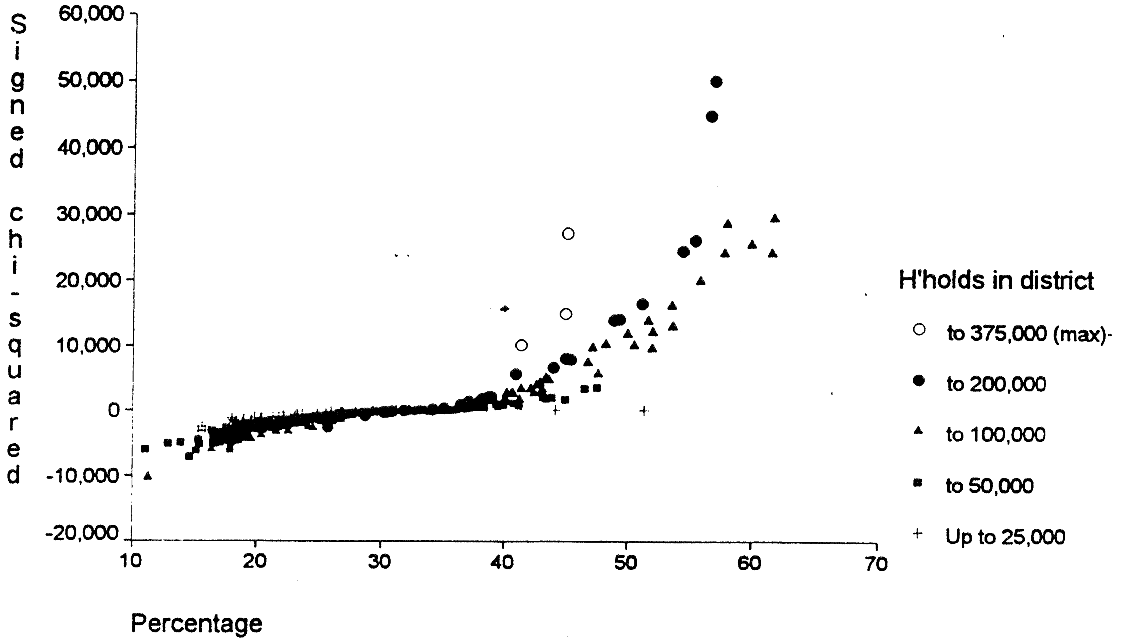

signed chi-squared: England 0; mean 475; SD 6,370; skewness 4.05

Percentage: England 32.4; mean 28.9; SD 10.4; skewness 0.93

Figure 2: Percentage and signed chi-squared statistics,

Residents in households lacking amenities

366 Districts of England

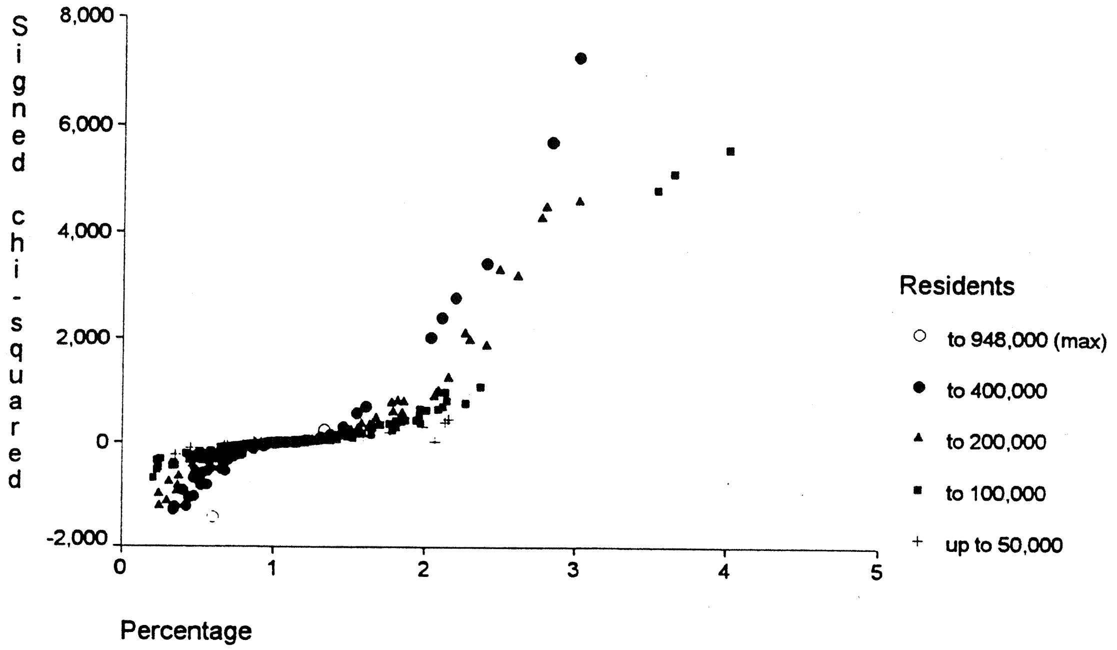

signed chi-squared: England 0; mean 134; SD 943; skewness 4.2

Percentage: England 1.09; mean 1.13; SD 0.58; skewness 1.34

- *The signed chi-squared identifies a greater number of needy individuals than does the percentage, within the highest scoring areas.*

Table 4 **The cumulative inclusion of needy individuals in areas identified as having high concentrations of need: signed chi-squared and percentage compared.**

Number of districts included	Percentage of households without a car Districts ranked by: %	χ^2_s	Percentage of residents lacking amenities Districts ranked by: %	χ^2_s
50	35.4	39.2	24.7	27.2
100	56.2	57.3	41.5	43.6
150	69.5	68.1	56.0	57.5
200	79.6	77.9	71.2	69.2
250	87.3	84.8	81.0	78.9
300	93.1	91.7	90.9	88.2
350	98.7	98.2	98.9	96.5
366	100.0	100.0	100.0	100.0

This is clearly a beneficial result when identifying areas of highly concentrated need for resource allocation. The advantage of the signed chi-squared over the percentage holds for any cut-off point above the reference area percentage. For example, the 50 districts with least car-ownership include 39.2% of all households in England without a car when the districts are identified using the signed chi-squared, while 35.2% are included when the districts are identified using the percentage. The graph and table are also a reminder that most programmes that target geographical areas miss the majority of those in need. They address inequality of need rather than the root causes of need.

These three welcome properties of the signed chi-squared statistic have a sting in the tail.

- *For areas with a percentage below that of the reference area, the chi-squared statistic behaves perversely in resource allocation.*

We have already seen in the case of Leeds and Great Grimsby for residents in households that lack basic amenities, that a large area below the reference area's percentage level of need will be pushed below smaller areas, even when those smaller areas may have a lesser need in both absolute and percentage terms. This is apparent from the construction of the chi-squared statistic, and in practice from Figures 1 and 2. Figure 3 shows the equally undesirable consequence that for areas with percentages below the reference area, a ranking based on the chi-squared statistic includes fewer of those in need than does a ranking based on percentages.

As identified among 366 English Districts when ranked by percentage (solid lines) and signed chi-squared (dotted lines)

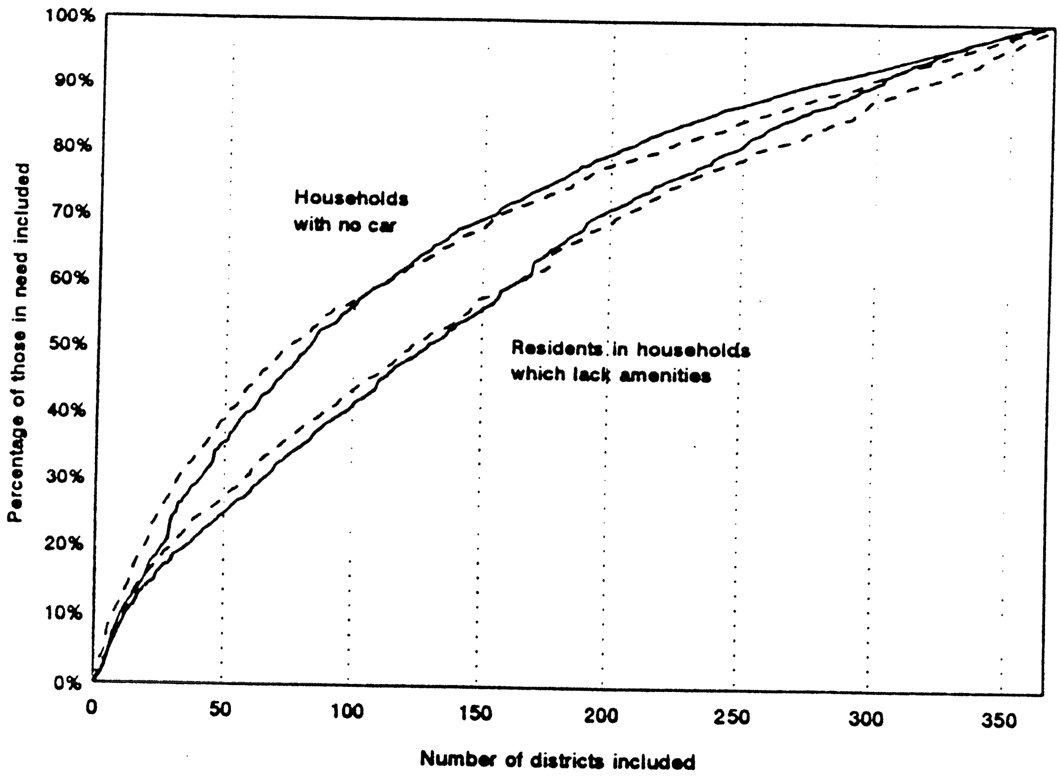

The signed chi-squared may serve to emphasise reliable differences from a national average for percentages based on small areas, a useful property for mapping. However it does not deal with heterogeneity appropriately for areas below the reference. The rank of large areas is lowered compared to small ones with a similar percentage need. For this reason alone, the signed chi-squared is not an appropriate basis for distribution of resources to areas of low relative need.

This problem is not so important if a signed chi-squared indicator is used to identify only high concentrations of need. However, some serious doubts might remain about an index combined of various such indicators. There may well be areas which are highly in need (high above the reference) on some indicators but below the reference on others - the two districts used as examples are cases in point on this issue too. In analysis of census Enumeration Districts, such cases would happen frequently. The down-weighting of large areas with a value below the reference area on a single indicator would unduly affect such areas' overall index score. Even when marginal, such effects can push an area into or out of the set of areas identified for priority resource allocation.

Finally, there is one further property of the signed chi-squared statistic that may be considered as a serious drawback.

- *Choice of reference area in the chi-squared statistic affects the ranking of areas.*

Figure 4 illustrates this property by showing the value of the signed chi-squared statistic for two hypothetical areas with percentages 38% and 40%, the first area being twice the size of the second. This is by no means an extreme case since local authority Districts and Enumeration Districts in England each vary by much more than a factor of two (in fact they by a factor of more than five, see Cole (1993)). The chi-squared value of the larger area exceeds that of the smaller area when the reference percentage π is low, but drops beneath it as π moves between 33% and 34%. In general large areas gain if a low reference π is chosen, and small areas gain from choice of a high reference π. The effect on chi-squared of size of area is most effective in distinguishing areas which are far from the reference π.

Figure 4: Signed chi-squared
The effect of varying the reference area percentage
on two areas of different size

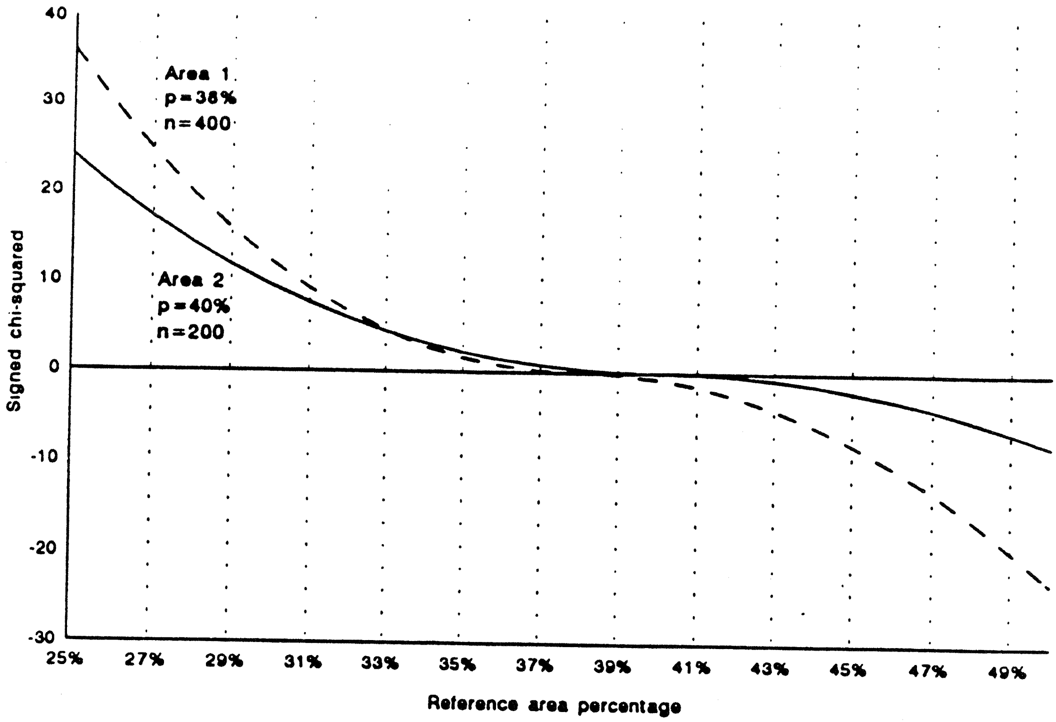

Table 3 provides a practical example of this effect. In the comparison of Leeds and Great Grimsby, choice of either of these districts as reference area (as might seem appropriate if the comparative study had been carried out on behalf of either of these districts' local authorities) would give Leeds considerably lower signed chi-squared value than Great Grimsby, rather than the slightly higher score where England is used as the reference.

In some applications, as where a national body is comparing all its constituent areas, a particular reference area may be clearly the most appropriate choice. But in general the choice of a reference area will depend on the perspective of the research application, and the use of the chi-squared statistic gives different results according to different perspectives.

5 Discussion

5.1 Comparison of count, percentage, and signed chi-squared

This paper has considered three alternative statistics when constructing a social index to distinguish geographical areas - the raw count of those affected, the percentage these make up of a relevant population, and the signed chi-squared statistic which combines the deviation of these two from some standard norm. It has identified two major purposes of distinguishing areas from one another in the context of resource allocation. Table 5 summarises the pros and cons of each statistic in relation to these two purposes.

Table 5: Summary of findings

	Distribution of resources by identifying aggregate need in each area:	*Geographic ring-fencing of resources, by identifying concentrations of need:*
Raw counts O, the observed number in need.	•Appropriate. Multiple indicators can be combined, retaining ease of interpretation.	•Makes no distinction between a thinly spread need and a concentrated need. Not suitable.
Percentages $p=100*O/n$	•Used as a per capita multiplier to a relevant population. •Can be misleading if the percentage is not derived from the relevant population.	•Unreliable for census statistics for small Enumeration Districts. •Significant concentrations within large populations are often missed.
Signed chi-squared $\chi^2_s=(p-\pi)^2n/\pi(100-\pi)$ Weights the percentage by the number in need, measured in relation to a reference percentage π.	•Not suitable if any area has a percentage below the reference.	•Small areas need a large percentage to be identified, and thus the unreliability of census small area statistics is acknowledged, as is larger areas' heterogeneity. •More needy individuals are contained within the highest scoring areas than when using percentages. •Treatment of low-value areas is not appropriate; although this is unimportant to identifying high-concentration areas, there may be an unwelcome effect within a combined index. •Choice of a reference area may be debateable, and affects the ranking of each area.

For distribution of resources to all or a chosen subset of areas, whether for funding schools, districts or regions, the percentage can provide a per capita amount for each area, used as a multiplier of some relevant population that the resources are aimed at. As an alternative, the simplicity and transparent role of raw counts have much to recommend them. The signed chi-squared is neither designed nor appropriate for this application, particularly if any low-need areas are involved.

For identification of areas with highest concentrations of need, the conclusions are rather different since the focus is not only on aggregate numbers in need but on high numbers in need within relatively small populations. Here the percentage is clearly not ideal because it ignores the concentrations that may be found within large areas and the unreliability of data for small areas: for both reasons the percentage sets larger areas at a disadvantage. But the jury must remain out if the signed chi-squared is considered as an alternative, because of its dependence on choice of comparator area and its perverse behaviour for areas below that comparator area. The signed chi-squared statistic may be appropriate in cases where a single reference norm is acceptable and only areas with the very highest concentrations are reported.

5.2 Transformations, standardisation and weighting when combining signed chi-squared variables.

Great care must be taken when combining more than one indicator into a social index, as both variance and skewness are exaggerated in the signed chi-squared statistic compared to the percentage (as can be seen from Figures 1 and 2). Standardisation of variables to equalise their variance (for example creating z-scores of signed chi-squared scores), and transformations to reduce the skewness of variables, may avoid these characteristics giving unexpected and undesired weight to one or other variable, just as with percentages (Senior, 1991). Complete normalisation of variables can be gained in common statistical packages after ranking a variable.

However, normalisation is only justified when used in order to allow statistical inference, which is not involved in the resource allocation discussed here. Paul Metcalf (1994) points out how the logarithmic transformation of chi-squared has made District differences close to the English average more influential than other District differences in the Index of Urban Conditions already referred to. David Martin, Martyn Senior and Huw Williams (1994) judge that no transformations are required when constructing an index of aggregate need. In general transformation loses information that may be important to policy, and adds an extra technical barrier to potential user's understanding of social indices, and therefore should be fully justified.

5.3 Alternatives measures of geographical concentrations of need

Are there alternative measures of concentration that address the unreliability of the percentage in small areas and its lack of attention to heterogeneity of large areas, but do not have the unfortunate properties of the signed chi-squared statistic for areas below the somewhat arbitrary reference percentage? Two alternatives that are adjustments of the chi-squared statistic are:

- Set the reference area to $\pi=0$, and express the statistic simply as p^2n or equivalently pxO, the product of the percentage and the number affected.

- Set the chi-squared statistic to 0 for all areas for which $p < \pi$, instead of changing the sign of such areas.

Both these alternatives retain the attractive weighting of percentage by size of area, but lose the negative aspects of negative values of the signed chi-squared.

A third alternative, used by Sally Holtermann (1975) for Enumeration Districts in Britain in 1971 and at Bradford Council (1993) for Enumeration Districts in one local authority area in 1991, is to:

- Provide a threshold for the population affected, so that areas that do not meet this threshold are considered unreliable or simply not large enough to warrant attention.

The former report excludes Enumeration Districts with less than 50 residents in households, while the latter resets high percentages to the District mean if less than 5 residents are affected. Both reports also transform variables to their ranks and combine variables not through a linear weighting function but through identifying those areas which are extreme on more than one variable: areas of multiple stress.

These three methods of identifying concentated need do seem to suffer, like the signed chi-squared, from arbitrariness in their construction. If heterogeneity of larger populations is considered a major issue, it might be fruitful to find a statistic which is directly related to what is known of the geographical heterogeneity of each variable, demanding of that statistic that it:

- Set a large area's concentration equal to a smaller area's concentration if there is an even probability that it contains an area of the smaller size but greater percentage.

The difficulty of finding such a statistic and of estimating the measures of heterogeneity that it demands serve to illustrate the difficulty of approaching this subject without recourse to arbitrary decisions even at this level of choice of basic unit.

When areas as large as wards or local authority districts are compared, the main drawback of the percentage is the heterogeneity of the larger units, which may have concentrations of need hidden within them. But then there may exist data for smaller areas (for example census enumeration districts or postcodes); this data should be used to search for concentrations directly and precisely. If the social policy application requires a minimum number of individuals in need in each target area, and/or a required minimum population, one could search for the areas with greatest percentage need, combining adjacent small areas until the size criteria are met.

Thus technological power may provide new solutions to the problem of identifying high social concentrations within a 'finely geographied' national census data set. If the size of area large enough for a relevant social programme is fixed, a geographical information systems (GIS) algorithm may be specified to:

- Scan all possible areas of specified size to identify those with the highest percentages, respecting local authority boundaries (or not, as the social programme may require).

Perhaps the search for one dimension of concentrations of social need is wrong. The strength of the Index of Urban Conditions referred to, is not so much its use of the chi-squared measure as its provision of a variety of indices, measuring different aspects of the geography of poverty. However, this paper will have served a useful purpose if it has shown that the choice of unit of measurement is important in indices of social need, and the success of that choice can be judged by the index having properties relevant to its use in resource allocation.

Appendix: Census variables defined.

No case is made for either of two variables used in this paper to be without fault as a measure of general need or poverty. They were chosen as examples because of their familiarity within the literature, and because they are not highly correlated to each other.

With reference to the census small area statistics (London Research Centre, 1993), households with no car have been extracted as cell S200131. To calculate a percentage it has been divided by the total number of households, S200001.

Residents in households lacking amenities are those in non-permanent accommodation (caravans etc) and those without exclusive inside toilet and bath/shower facilities. It has been extracted as the sum of cells S200181 and S200149. To calculate a percentage it has been divided by the total residents in households, the sum of cells S200141 and S200149.

Acknowledgements:

All census data is Crown Copyright. Comments on a draft were helpfully made by Keith Cole, Angela Dale, Ed Fieldhouse, Richard Henderson, Helen Quigley, and Martyn Senior.

References:

Bradford, Michael, Brian Robson, and Rachel Tye (1994) 'Constructing an urban deprivation index'. Environment and Planning Series A (forthcoming).

Bradford Metropolitan District Council (1993b) 'Areas of stress within Bradford District: a report from the 1991 census and other sources'.

Browne, Martin (1992) 'Home care area budget allocation. Paper 1: Testing the alternatives'. Bradford: City Council Social Services Policy and Information Unit.

Chipping, Hilary (1994) 'Data for allocation of resources', in 'Regional and local statistics', proceedings of the Statistics User's Council Annual Conference 16th November 1993. Esher: IMAC research.

Cole, Keith (1993) 'The 1991 Local Base and Small Area Statistics', in 'The 1991 Census User's Guide', eds Angela Dale and Catherine Marsh, 1993. London: HMSO.

CRU/OPCS/GRO(S) (1980) 'People in Britain - a census atlas'. London: HMSO.

Davies, Hugh, Heather Joshi, and Lynda Clarke (1995) 'Is it cash the deprived are short of?', Journal of the Royal Statistical Society series A, forthcoming.

Derbyshire, Mike (1993) 'Changes in local government finance'. Journal of the Royal Statistical Society series A 156(3): 333-335.

DoE (1983) 'Urban Deprivation' Information Note 2, Inner Cities Directorate, 2 Marsham Street. London: Department of Environment.

Government Offices for the Regions (1994) 'Bidding guidance: a guide to funding under the single regeneration budget'. London: HMSO.

Holtermann, Sally (1975) 'Areas of urban deprivation in Great Britain: an analysis of 1971 census data'. Social Trends 6: 33-47.

Jenkins, John (1994) 'Using the 1991 census for the review of childcare', in 'Research for policy 1994' proceedings of the Local Authorities Research and Intelligence Association conference 28-30 March. LARIA: Newcastle.

London Research Centre (1992) 'SASPAC User Manual Part 2'. Local Government Management Board and London Research Centre: London.

Martin, Huw, Martyn Senior and Huw Williams (1994), 'On measures of deprivation and the spatial allocation of resources for primary health care', Environment and Planning series A, forthcoming.

Metcalf, Paul (1994) 'Comparing parish or district deprivations using the chi-squared score'. Unpublished paper: 67 Park Hall Road, Walsall WS5 3HL.

Rhind, David (ed.) (1983) 'A census user's handbook'. London: Methuen.

Senior, Martyn (1991) 'Deprivation payments to GPs: not what the doctor ordered'. Environment and Planning C, 9: 79-94.

Senior, Martyn (1994) 'The English Standard Spending Assessment system: an assessment of the methodology'. Environment and Planning C: Government and Policy, 12: 23-51.

Simpson, Steve (ed.) (1993) 'Census indicators of local poverty and deprivation: methodological issues'. Proceedings of a day workshop organised by the Manchester Census Group and Local Authority Research and Intelligence Association. London: LARIA.

Thunhurst, Colin (1985) 'The analysis of small area statistics and planning for health'. The Statistician, 34: 93-106.

Other titles available from the Census Microdata Unit:

Census Microdata Unit Occasional Papers

No 1. *Problems of Imputation in the 1991 Census,* (1 899005 02 1), Amarjit Sandhu

No 2. *Bias , Sampling Error and Coverage: the preliminary validation of the Samples of Anonymised Records from the 1991 Census,* (1899005 03 X), Steve Simpson, Ed Fieldhouse & Amarjit Sandhu

No 3. *An Introductory Guide to Analysing the SARs,* (1899005 04 8), Liz Middleton

SARs User Guide (1 899005 01 3)

Manchester Census Group

The Ethnic Dimensions of the 1991 Census: A Preliminary Report, (1 89900500 00 5), Roger Ballard & Virinder Singh Kalra

£3.00

ISBN 1 899005 05 6

Further copies of this paper may be obtained from:

Census Microdata Unit
Faculty of Economic and Social Studies
The University of Manchester
Oxford Road, Manchester, M13 9PL