Review of Economic Studies (1982) XLIX, 273–290 0034-6527/82/00210273$00.50

The Estimation and Interpolation of Inequality Measures

FRANK A. COWELL

and

FATEMEH MEHTA

London School of Economics

Alternative methods of computing estimates of inequality measures from grouped data are critically examined in terms of their theoretical and empirical properties. The use of a simple "split-histogram" technique of interpolation is explained and supported. Theoretical and empirical support is also provided for the "$\frac{1}{3}/\frac{2}{3}$ rule"—a simple computational procedure for a point estimate of an inequality measure derived from its standard grouping bounds.

1. INTRODUCTION

If one has to work with published data on income distributions can one reliably estimate measures of inequality? Not surprisingly the answer to this depends on the degree of detail to which the information is presented. More interestingly it turns out that good estimates can be provided with comparatively limited information and that simple computational methods are available that prove to be robust under a wide range of circumstances. In this respect our results have strong implications for the collection and presentation of official statistics.

We shall concentrate in the main on the extremely common case where the data are presented in grouped form and one knows, for a given set of income intervals, the number of people in each interval and the mean income of each interval, but nothing more. Using only this information it is well known that one may calculate upper and lower bounds within which the "true" value of the inequality measure must lie,[1] so if these bounds are sufficiently narrow then obviously we obtain reliable estimates immediately. However, there are many circumstances under which the bounds are such that the range of possible values of the statistic is quite broad, so that the question of a suitable intermediate value immediately arises. As it happens the simplest guess—midway between the bounds—is always a bad one. However, we shall show that a guess that is very nearly as simple turns out usually to be a very good one.

It is fairly evident that to arrive at such an intermediate value by a method more persuasive than guessswork one must make some assumption about the way in which the population is distributed within the income ranges with which we are confronted. Part of our task, therefore, is to consider what criteria should be satisfied by a suitable method of interpolating on the distribution. For example, does it matter whether or not the assumed frequency distribution exhibits continuity? How sensitive is the intermediate, interpolated value of the inequality measure to alternative assumptions about the underlying distribution? How good are very simple interpolation methods? In appraising the various possible interpolation methods we shall insist on only one absolute requirement—that the implicit frequency function shall be statistically feasible which in practice means that the frequency function must be able to yield the observed interval means and population densities without ever requiring (locally) a negative frequency.[2]

It is also evident that the statistical properties of an inequality measure may depend on its theoretical characteristics. Accordingly we shall examine which measures are suitable in principle for estimation from such data, how the measures themselves can provide a test criterion for alternative interpolation methods, and how the reliability of particular estimates depends on the theoretical sensitivity of the inequality measure to transfers in different income ranges.

2. INTERPOLATION: THE GENERAL PROBLEM

We shall assume that the income distribution data are a set of grouped observations of points in $[0, \infty)$. Specifically we shall suppose that there are ω income classes

$$[a_1, a_2), [a_2, a_3), \ldots, [a_\omega, a_{\omega+1})$$

where $0 \leqq a_1 < a_2 < \cdots < a_\omega < a_{\omega+1} \leqq \infty$. Apart from the income interval boundaries a_θ, $\theta = 1, \ldots, \omega$ we shall assume that we have observations of n_θ, the numbers of people with incomes in $[a_\theta, a_{\theta+1})$ and of μ_θ, the mean income of people with incomes in that interval.

Since we do not know the true density function $\hat{f}(y)$ of income y it is permissible to assign people arbitrarily to particular incomes subject only to two key constraints; (a) for any arbitrary assignment the number in interval θ equals n_θ, and (b) each interval mean in the assignment equals μ_θ. In particular, given any inequality measure I that is an S-convex function from the space of incomes to the real line, we can find two assignments of the population that provide, respectively, a minimum and a least upper bound of I subject *only* to these constraints namely

$$\left.\begin{aligned} f_1(y) &= \frac{n_\theta}{n} \quad \text{if } y = \mu_\theta, \quad \theta = 1, 2, \ldots, \omega \\ &= 0, \quad \text{otherwise} \end{aligned}\right\} \tag{1}$$

$$\left.\begin{aligned} f_2(y) &= \lambda_\theta \frac{n_\theta}{n} \qquad \text{if } y = a_\theta \\ &= [1-\lambda_\theta]\frac{n_\theta}{n}, \text{ if } y = a_{\theta+1} \\ &= 0, \qquad \text{otherwise} \end{aligned}\right\} \tag{2}$$

where $n \equiv \sum_{\theta=1}^{\omega} n_\theta$ and $\lambda_\theta \equiv [a_{\theta+1} - \mu_\theta]/[a_{\theta+1} - a_\theta]$.

If one has the *additional* information that within a given interval θ, $\hat{f}(y)$ is a decreasing function one may improve on these two assignments by using, respectively,

$$\left.\begin{aligned} f_3(y) &= \frac{n_\theta}{2n[\mu_\theta - a_\theta]}, \quad \text{if } a_\theta < y < 2\mu_\theta - a_\theta, \\ &= 0, \qquad \text{otherwise,} \end{aligned}\right\} \tag{3}$$

and

$$\left.\begin{aligned} f_4(y) &= \frac{n_\theta}{n}\frac{a_\theta + a_{\theta+1} - 2\mu_\theta}{a_{\theta+1} - a_\theta}, \quad \text{if } y = a_\theta, \\ &= \frac{2n_\theta}{n}\frac{\mu_\theta - a_\theta}{[a_{\theta+1} - a_\theta]^2}, \qquad \text{elsewhere.} \end{aligned}\right\} \tag{4}$$

If in a given interval the decreasing frequency assumption is *not* true then let $f_3(y) = f_1(y)$ and $f_4(y) = f_1(y)$ over that interval. Writing $I_1, \ldots, I_4, \hat{I}$ for the values of I corresponding to $f_1(y), \ldots, f_4(y), \hat{f}(y)$ respectively it may be shown that whatever the underlying distribution

$$I_1 \leqq I_3 \leqq \hat{I} \leqq I_4 \leqq I_2.$$

Of course if the gap $I_4 - I_3$ or the gap $I_2 - I_1$ is sufficiently small then the problem of estimation from this kind of data is virtually solved. However, we shall see in Section 6 that $I_4 - I_3$ may be significantly large so that we have the problem of finding an intermediate value for I. To do this we need to find some "plausible" hypothetical frequency function $f(y)$ so that the computed value of I lies within the required limits. We shall refer to this as the "interpolation problem". Clearly the method of interpolation may significantly affect the computed values of I, and much may depend on what we mean by a "plausible" $f(y)$. Let us consider ten properties which it might be desirable that the hypothetical interpolated distribution should possess.

Properties:

1. Constraints (a) and (b) are satisfied.
2. $f(y) \geqq 0$.
3. Continuity of $f(y)$ within any interval θ.
4. First-order continuity at a_θ, $\theta = 2, 3, \ldots, \omega$.
5. Second-order continuity at a_θ, $\theta = 2, 3, \ldots, \omega$.
6. $\lim_{y \to a_{\omega+1}} f(y) = 0$.
7. $\lim_{y \to a_{\omega+1}} f'(y) = 0$.
8. "Few" turning points of $f(y)$.
9. A "small" range of $f(y)$ for any given interval.
10. Closed form integration of inequality measures.

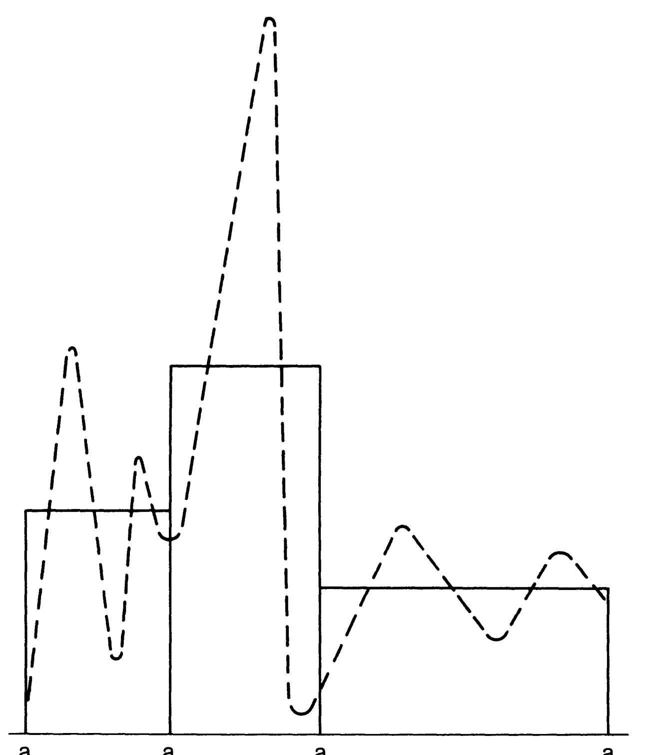

FIGURE 1

A pathological interpolation function

Apart from 1 and 2, all of these are tentative and some are very vague. What we shall demonstrate, however, is that this vagueness does not matter much in practice. On occasion we shall weaken 3 to piecewise continuity, rather than continuity. Properties 8 and 9 will be interpreted below. The idea is to avoid wild fluctuations in the imputed frequency distribution. The reason for this is easy to see in Figure 1, where the observed frequencies are drawn in as a histogram. The arbitrarily-imposed income interval bounds give us no reason to suppose the existence of substantial peaks and troughs; had the official statistical source redrawn the bounds, it is quite likely that the shape of the observed distribution would have been broadly similar. Hence it is desirable to avoid an interpolated curve of the form illustrated, though of course it is difficult to give a neat, tractable and general form to the kind of mathematical restriction that avoids this problem.

3. METHODS OF INTERPOLATION

To the best of our knowledge there is no single, simple function which will satisfy all ten points in general. However there are a number of possible compromises which we examine in this section.

Piecewise Paretian interpolation. The so-called Pareto distribution of type I has often been suggested as a general functional form for the upper tail of income distribution. So it is natural to consider a Paretian density function for interpolation within intervals, thus:

$$f_5(y) = A_\theta y^{-\alpha_\theta - 1}, \qquad y \in [a_\theta, a_{\theta+1}) \tag{5}$$

where A_θ and α_θ are parameters chosen to ensure that conditions (a) and (b) are satisfied. Clearly $f_5(y)$ has Properties 1–3 and 8, although there is no reason to suppose that 4, 5 and 9 will be satisfied.

Properties 6 and 7 can be guaranteed if $a_{\omega+1} = \infty$. We shall find that Property 10 presents no problem.

Polynomial interpolation. Consider first the use of a polynomial spline using ω functions of the form

$$f_6(y) = \textstyle\sum_{k=0}^{K} \gamma_{\theta k} y^k, \qquad y \in [a_\theta, a_{\theta+1}). \tag{6}$$

This form can satisfy all of Properties 1 to 7 for $K \geqq 3$. Property 10 is in practice easily satisfied for polynomial of any order given the inequality measures we shall describe below. However in practice there may be considerable problems with Properties 8 and 9, since of course there can be up to $\omega[K-1]$ turning points in the resulting spline curve. Because of the possibility that this smooth curve may exhibit rather implausible "cork-screwing", we were motivated to examine some rather simpler forms. Perhaps the most obvious is to consider polynomial forms with $K < 3$, namely the straight line and the quadratic. The advantage of the straight line

$$f_7(y) = \gamma_{\theta 0} + \gamma_{\theta 1} y, \qquad y \in [a_\theta, a_{\theta+1}) \tag{7}$$

is of course that it is extremely easy to calculate and to use. The disadvantage is that it is quite likely that once $\gamma_{\theta 0}$ and $\gamma_{\theta 1}$ are found so as to ensure that Property 1 holds, $f_7(y)$ may intersect the horizontal axis, thus violating Property 2. A similar problem exists with the quadratic and some other polynomials, and we discuss this practical problem in Appendix A.

A number of authors[4] have used a cubic to approximate sections of the Lorenz curve. The form of the frequency function in each interval is then

$$f(y) = C_{\theta 0}[C_{\theta 1} + y]^{1/2}.$$

This is in fact a special case of the Pareto type II distribution, and we may expect its performance in practice to be rather like a type I curve, $f_5(y)$. Moreover the computation of some inequality measures is rather tedious using this method. For these two reasons we shall not consider this further.

Split-histogram interpolation. This apparently unusual title refers to an extremely useful and simple general method of interpolation. Let

$$\left.\begin{aligned} f_8(y) &= \frac{n_\theta}{n} \frac{a_{\theta+1} - 2\mu_\theta + b_\theta}{[a_{\theta+1} - a_\theta][b_\theta - a_\theta]}, \qquad y \in [a_\theta, b_\theta) \\ &= \frac{n_\theta}{n} \frac{2\mu_\theta - b_\theta - a_\theta}{[a_{\theta+1} - a_\theta][a_{\theta+1} - b_\theta]}, \qquad y \in [b_\theta, a_{\theta+1}) \end{aligned}\right\} \tag{8}$$

where b_θ is an arbitrary point in $[a_\theta, a_{\theta+1}]$. Properties 1, 2, 8, 9, 10 are satisfied, for some suitable choice of $b_1, \ldots, b_\omega$.[5] However, the function obviously does not provide a neat upper tail for finite $a_{\omega+1}$, and it will in general be discontinuous at $2\omega - 1$ points in the income range. The function is illustrated in Figure 2. Note that it is essential in general to split the histogram for this kind of problem since the simple histogram:

$$f_9(y) = \frac{n_\theta}{n} \Big/ [a_{\theta+1} - a_\theta], \qquad y \in [a_\theta, a_{\theta+1}) \tag{9}$$

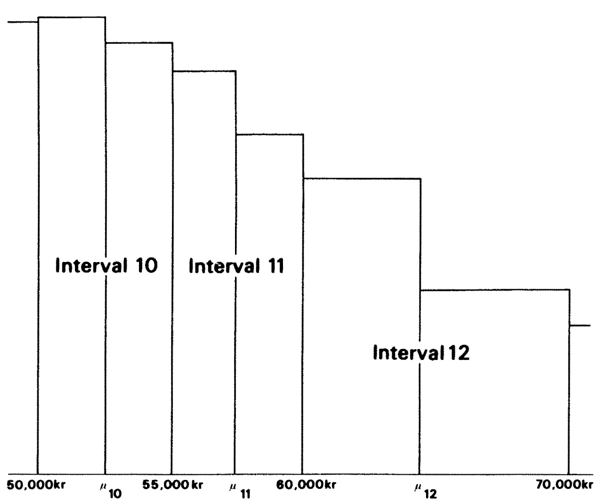

FIGURE 2
The split histogram. Intervals 10 to 12 for Sweden 1977

will be inconsistent with the condition that the mean in each interval shall equal μ_θ (see Property 1).

There is no reason of course why more than one method of interpolation should not be used—indeed we could use a different formula for each interval. With two exceptions, though, there seems to be little point in doing this and we might as well find an acceptable compromise method and stick to this throughout the income range. One may always use a different compromise for different types of data sets. The first of the two exceptions is where a generally useful compromise breaks down in a special case. An example of this is where a generally satisfactory polynomial spline is found to cross the axis within some interval: within that interval one may then wish to use an alternative cruder interpolation function, whilst retaining the general formula elsewhere.[6] The second exception relates to the top interval, in most cases official statistical sources leave $a_{\omega+1}$ unspecified. One may therefore *either* assume some arbitrary upper limit $a_{\omega+1}=L$ such that $f(y)=0$ if $y>L$, *or* let $a_{\omega+1}=\infty$, neither of which is entirely satisfactory. In the first case the results on I_2, I_4 and the various interpolations may be rather sensitive to the particular choice of L—although one can in practice place reasonable limits upon this choice.[7] In the second case, an obvious but rather arbitrary procedure is to assume that the density tends to zero in a specific fashion—that is one chooses a function where $f(y)\leqq 0$ and where Properties 6 and 7 hold asymptotically at infinity, for example a Pareto tail. Unfortunately, the required estimate of α may be quite sensitive to the grouping in the upper tail, and even if the distribution over the last few intervals is approximately Paretian, it may not be legitimate to extrapolate the functional form to infinity.[8]

4. THE D-MEASURE

We are left, then, with a number of more or less satisfactory methods of interpolating a density function and hence of making single-index inequality comparisons, rather than merely reporting and comparing grouping bounds. But how much difference will the interpolation methods make in practice? Will one interpolated function lie rather close to another? Can we find a simple expedient compromise procedure that is almost as good as any other?

To answer these questions we need first to examine inequality measures themselves. An inequality measure is an S-convex function from the space of incomes to the real line. We are especially interested in the subclass of inequality measures that are decomposable by population subgroups where the subgroups are non-overlapping—known as Non-Overlapping—Decomposable Inequality (NODI) measures. If further we insist that such measures are scale-independent so that doubling or halving all incomes leaves measured inequality unaltered, then we are restricted to either the Gini index:[9]

$$I^G=\frac{1}{2\mu}\int_0^\infty\int_0^\infty |y-z|f(y)f(z)dydz \tag{10}$$

or the Generalized Entropy class of inequality measures:[10]

$$I^\beta=\frac{1}{\beta[\beta+1]}\left[\int_0^\infty\left[\frac{y}{\mu}\right]^{\beta+1}f(y)dy-1\right] \tag{11}$$

where β is a parameter which may take any real value. Given measures of this form we may then write:

$$I^G=\sum_{\theta=1}^{\omega}\frac{n_\theta^2\mu_\theta}{n^2\mu}I^{G\theta}+\frac{1}{2\mu n^2}\sum_{\theta=1}^{\omega}\sum_{\tau=1}^{\omega}|\mu_\theta-\mu_\tau|n_\theta n_\tau \tag{12}$$

$$I^\beta=\sum_{\theta=1}^{\omega}\frac{n_\theta}{\mu}\left[\frac{\mu_\theta}{\mu}\right]^{\beta+1}I^{\beta\theta}+\frac{1}{\beta^2+\beta}\sum_{\theta=1}^{\omega}\left[\left[\frac{\mu_\theta}{\mu}-1\right]\frac{n_\theta}{n}\right] \tag{13}$$

where $I^{G\theta}$, $I^{\beta\theta}$ are the values of the inequality statistic *within* interval θ. The second term in (12) and (13) is simply the between-interval inequality in each case, namely I_1.

Now consider the expressions

$$D_G\equiv\frac{I^G-I_1^G}{I_2^G-I_1^G},\qquad D_\beta\equiv\frac{I^\beta-I_1^\beta}{I_2^\beta-I_1^\beta}. \tag{14}$$

These ratios provide suitably normalized measures of the departure of the actual distribution from the minimum inequality f_1-distribution. If there is maximum within-interval inequality (the f_2-distribution) then $D_G = D_\beta = 1$. Moreover any given change in one of the within-interval inequalities, then this produces a proportionate increase in the D-measure.

Clearly this D-measure is a useful tool for appraising the closeness of any distribution to either of the two extremes f_1, f_2; hence we may also use this to measure the closeness of two arbitrary interpolations f^* and f^{**} that are bounded by f_1 and f_2.[11] In the absence of further information we may use $D^* - D^{**}$ to judge whether or not the two interpolation methods yield results that are distinguishable. For small perturbations in the underlying frequency function, this method may also be extended to inequality measures that are monotonic transformations of (10) and (11). Suppose, for example we consider $J \equiv \phi(I^\beta)$, where $0 < \phi' < \infty$, and two interpolations f^* and f^{**}. Using an obvious notation we have

$$D_J^* - D_J^{**} = \frac{\phi(I^{\beta *}) - \phi(I^{\beta **})}{\phi(I_2^\beta) - \phi(I_1^\beta)} \tag{15}$$

$$\simeq \left[\phi'(I^{\beta *})\frac{I_2^\beta - I_1^\beta}{J_2 - J_1}\right][D_\beta^* - D_\beta^{**}] \tag{16}$$

where the approximation involved in going from (15) to (16) is legitimate if f^* and f^{**} are sufficiently close such that ϕ' is locally approximately constant. Under these circumstances the first bracket in (16) is constant and $D_J^* - D_J^{**}$ is proportional to $D_\beta^* - D_\beta^{**}$.

This procedure has a simple intuitive appeal. Any statistical measure of "goodness-of-fit" involves some sort of distance function, and the choice of approximate measure will depend, in part, on the choice of distance function. Since we are ultimately concerned with the income distribution as a guide to inequality we have let our choice of distance function in constructing a measure of "closeness of approximation" be governed by the principles on which we construct our standard of inequality. "Relevant" income differences for the purposes of inequality measurement are reflected as "relevant" differences between two proposed frequency functions.

Of course the D-measures are not unique, because in general inequality measures do not provide unique rankings. It can happen that the D-measures associated with different base NODI measures reveal different results because certain NODI measures will be particularly sensitive to perturbations in the underlying f in specific parts of the distribution. To see this, consider an arbitrary perturbation of f within interval θ. To be admissible conditions (a) and (b) in Section 2 must still hold, so that the perturbation can be expressed as the sum of elementary transfers of the sort: "decrease the income of a man with \$$y$ by an amount \$$dy$ and increase the income of a man with \$$y+\Delta$ by \$$dy$". The resulting variation in D for any one such elementary transfer is given by

$$dD = \frac{1}{I_2^\beta - I_1^\beta}\left[\frac{1}{\beta}\left[\frac{y+\Delta}{\mu}\right]^\beta - \frac{1}{\beta}\left[\frac{y}{\mu}\right]^\beta\right]dy \tag{17}$$

and the total effect on D is found by integrating (17) for the entire variation in f.

Evidently the size of the total change in D for some perturbation in f will depend on three things: the value of β, the distribution of the population within and amongst the intervals, and the interval widths. The underlying distribution is important because of the weighting effect in constructing the total D. The interval width, of course, places a constraint on the size of Δ we need to consider in (17). Now take the parameter β. Differentiation of (17) with respect to β reveals that for given Δ, if $y > \mu$, $y + \Delta > \mu$ and $dy > 0$, then dD increases with β if $\Delta > 0$ and decreases with β if $\Delta < 0$. On the other hand if $y < \mu$, $y + \Delta < \mu$ and $dy > 0$, dD decreases with β if $\Delta > 0$ and increases with β if $\Delta < 0$. Hence in one of the upper intervals—where we may safely assume $\mu < a_\theta < y$ and $\mu < a_\theta < y + \Delta$—the absolute size of the change in D increases with β, whereas in the bottom intervals—where $\mu > a_{\theta+1} > y$ and $\mu > a_{\theta+1} > y + \Delta$—the absolute size of the change in D decreases with β. Hence we expect distance based on NODI measures with relatively high β-values (such as the variance, with $\beta = 1$, or Theil's index with $\beta = 0$) to be relatively sensitive to perturbations in the top interval(s), and distance based on NODI measures with low β-values (such as Atkinson's index with inequality aversion greater than unity—$\beta < -1$) to be sensitive to perturbations in the bottom intervals.

However, although the D-measure *may* in general be sensitive to the specific assumptions cited above, in practice the regularity of observed frequency distributions may mean that the computed D-measures tend to a particular uniform value. In fact for some analytically tractable functional forms of the frequency distribution an amazingly simple and attractive rule for the value of the D-measure can be established: for the Gini coefficient D_G is two-thirds—for every other NODI measure D_β is (approximately) one third.[12] However, whether this rule is more generally applicable can of course only be resolved by empirical investigation, which we undertake in the next two sections.

5. EMPIRICAL RESULTS: MICRODATA

In order to do a comparison of the behaviour of the various methods of interpolation that were discussed in the previous section, as opposed to simply estimating the bounds for the inequality measures, we first turned to microdata. Unfortunately, the large data sources which are used by official statistical bureaux are not usually published and therefore this kind of information is usually not available to us.[13] To examine therefore, the behaviour of these various methods we are forced to use sample survey evidence. The data we have used for this purpose were drawn from the Michigan Panel Study of Income Dynamics and relates to total family income for 6003 families in 1968.[14]

The precise questions we wanted to investigate are these. Firstly, do the interpolation methods from an arbitrary grouping yield results that are good approximations to results one would have by directly using the microdata? Secondly, is the "one third/two thirds" rule supported for D-measures from such arbitrary groupings? To answer these questions we obviously need a little preliminary clarification. As far as the choice of some arbitrary grouping of the data goes we decided to use simply the same grouping of incomes as that used by the Current Population Survey in its published tables. For the purpose of comparison of the inequality measures obtained from the raw data, with those obtained from the grouped data in the next section we concentrated on the two methods of interpolation 5 and 8. Our criterion of judgement was based on a confidence interval for each inequality measure obtained from the raw data: we examined whether the interpolated values from f_5 and f_8 did in fact fall within that interval.[15]

The D-measures we used were all based on NODI measures (see equations 10 and 11). The particular cases of these that we used were the Gini coefficient, I^G, the coefficient of variation, $[2I^2]^{1/2}$, Theil's measure I^0 and Atkinson's measure with inequality aversion parameter $\varepsilon \equiv -\beta$, namely

$$1-\frac{1}{\mu}\left[\int_0^\infty y^{1-\varepsilon}f(y)dy\right]^{1/(1-\varepsilon)} = 1-[(\beta+\beta^2)I^\beta+1]^{1/(1+\beta)}.$$

In fact we used these particular cardinalizations of the measures which are familiar in the literature, involving non-linear transformation of I^β.[16] This leads to a slight bias in the D-measures as explained in Section 4, but as the reader can readily confirm, it in no way alters the substantive results.

The results related to this section appear in Table I, where the numbers in parentheses are the D-measures. Clearly the interpolated values lie very close to the values obtained from the raw data. Based on the criterion explained above they would be accepted as not significantly different from the directly computed values. The interpolated values do lie in the 95% confidence intervals and also the "one third/two thirds" rule holds pretty well.

TABLE I

Inequality measures for the Michigan microdata

	Lower bound	Interpolation method 5	Interpolation method 8	Upper bound	From raw data	95% confidence interval	
Gini	0·376117	0·377808 (0·662)	0·377821 (0·666)	0·378673	0·377608	0·361995	0·393220
C. of V.	0·754976	0·762294 (0·268)	0·764170 (0·337)	0·782234	0·762358	0·748705	0·776011
Theil	0·236886	0·239507 (0·296)	0·239828 (0·332)	0·245749	0·239345	0·209804	0·268886
Atkinson							
0·5	0·115183	0·116788 (0·290)	0·116948 (0·319)	0·120709	0·116611	0·100279	0·132794
1·0	0·223500	0·228343 (0·273)	0·228890 (0·304)	0·241209	0·227751	0·212958	0·242265
3·0	0·542239	0·590611 (0·303)	0·599016 (0·356)	0·701735	0·585075	0·570743	0·598061

Note: For the grouped data there are 21 income classes. For both methods minimum income is $500.

However the 95% confidence intervals—set up for each inequality measure are fairly wide. This reflects the fact that the estimated standard errors are appreciably large, and for this kind of significance test to be meaningful at all, these standard errors would have to be reduced by at least a factor of $\frac{1}{10}$, which in turn, implies that we must increase the size of the sample by a factor of 100. This limitation shows clearly the almost impossible task, at this stage anyway, of getting meaningful results by using microdata alone.[18] The next best choice—which provides very reliable results—is therefore to look at data which have been grouped already, and where the source is such that problems of large standard errors will not be present.

6. EMPIRICAL RESULTS: GROUPED DATA

Since the available sample survey data do not provide the basis for a definitive test of the track record of different methods of interpolation and of the "one third/two thirds" rule, we turned to income distribution data from tax returns which, though they have limitations in terms of their economic interpretation, are at least free of the particular statistical problems that arose in the previous section. These data,[19] for income before tax in Sweden, 1977, are presented in grouped form and the information made available to us for each interval is the number of people in each interval and the total income received in that interval, from which we can directly calculate the mean income of that interval. Since we are now working under this slight handicap, a number of issues of special interest arise.

The first point of interest is the comparison of the "Crude" bounds obtained from methods 1 and 2 and the "refined" bounds obtained from methods 3 and 4. Gastwirth (1972, 1975) recommended the use of the "refined" bounds as an accurate estimation procedure in preference to the use of "crude" grouping bounds, using as specific examples I^G, and I^β with $\beta = 0$ (Theil's inequality measure). Are these in fact sufficiently accurate? Obviously the strength of this method rests on the reduction in grouping error achieved by the refinement, and as Table II shows, this depends quite strongly on the choice of inequality measure, and in most cases is not particularly great (about 0·001). As inequality aversion increases so the 'refinement-reduction' diminishes from 43% for $\varepsilon = 0{\cdot}001$ to less than 1% for $\varepsilon = 10$. For this reason we cannot rely universally on the refined bounds

TABLE II

Inequality measure	Bounds 1	3	4	2
Gini	0·339419	0·340212	0·341571	0·341936
C. of V.	0·705777	0·711199	0·732457	0·752516
Theil	0·198498	0·199334	0·202692	0·204876
Atkinson				
$\varepsilon =$ 0·001	0·000198	0·000199	0·000203	0·000205
0·1	0·019743	0·019820	0·020149	0·020344
0·2	0·039301	0·039442	0·040103	0·040452
0·5	0·097170	0·097458	0·099293	0·099957
1·0	0·191634	0·192070	0·197261	0·198199
1.5	0·282570	0·283076	0·294389	0·295423
2·0	0·366685	0·367187	0·387690	0·388658
2·5	0·440494	0·440932	0·472604	0·473379
3·0	0·502211	0·502555	0·545398	0·545934
3·5	0·552147	0·552397	0·604748	0·605077
5·0	0·648831	0·648901	0·716458	0·716508
8·0	0·729228	0·729232	0·799525	0·799526
10·0	0·752706	0·752707	0·821331	0·821331

Notes:
1. Column Headings 1 and 2 refer to the Crude Lower and Upper Bounds.
2. Column Headings 3 and 4 refer to the Tight Lower and Upper Bounds.

as accurate estimates, and furthermore we have been content to compare our interpolated inequality measure with the crude bounds I_1 and I_2.

So, secondly, let us look at the results obtained from using the various interpolation methods 5 through 9. Will these different methods "agree", so that we can be confident about the point estimate of the underlying inequality statistic—even though we do not have details about the distribution within intervals? Examine Table III—our "standard" case where the number of income classes is 19, the lowest income class having been omitted, and we assume the top income class to have an upper bound of 2,000,000 kr.[20]

TABLE III

Inequality measures for Swedish grouped data (standard case)

	Lower bound	Interpolation method 5	6	7	8	9	Upper bound
Gini	0·339419	0·341095 (0·666)	0·341085 (0·662)	0·341088 (0·663)	0·341097 (0·667)	0·336965 (−0·975)	0·341936
C. of V.	0·705777	0·719601 (0·296)	0·721753 (0·342)	0·721602 (0·338)	0·721693 (0·340)	0·819006 (0·423)	0·752516
Theil	0·198498	0·200509 (0·315)	0·200631 (0·334)	0·200620 (0·333)	0·200649 (0·337)	0·219065 (3·225)	0·204876
Atkinson							
0·5	0·097170	0·098076 (0·325)	0·098074 (0·324)	0·098090 (0·330)	0·098101 (0·334)	0·103934 (2·427)	0·099957
1·0	0·191634	0·193785 (0·328)	0·193708 (0·316)	0·193786 (0·328)	0·193807 (0·331)	0·201783 (1·546)	0·198199
3·0	0·502211	0·517145 (0·342)	0·516643 (0·330)	0·517114 (0·341)	0·517173 (0·342)	0·524327 (0·506)	0·545934

Note: The original 20 income classes has been truncated so as to omit the first one. The Top income class is assumed closed with upper bound equal to 2,000,000 kr.

The performance of all of the interpolation methods (with the exception of 9 which is discussed later) is quite remarkable. The interpolated value appears to be insensitive to the widely differing methods of interpolation. Examining the D-measures, given in parentheses, there is indeed a clear tendency for these to approximate to $\frac{2}{3}$ in the case of the Gini coefficient and to $\frac{1}{3}$ in the case of the other measures, confirming the results that are available in analytically tractable cases—see Appendix A.

The next experiment was to test the robustness of these various measures if the presentation of the data had been less informative. For this purpose we originally merged the income classes together, keeping the first and the last intervals intact, so as to halve the total number of income classes to a mere ten. The results we obtained were quite strikingly similar to Table III. Hence we decided to go even further in merging intervals and progressively reduced the number of income classes to only five—see Table IV. A comparison with Table III shows quite obviously the robustness of the interpolated

TABLE IV

Inequality measures for Swedish grouped data: merged intervals

	Lower bound	5	6*	7*	8	9	Upper bound
Gini	0·307397	0·342332 (0·662)	0·340982 (0·637)	0·342539 (0·666)	0·342570 (0·667)	0·296988 (−0·197)	0·360157
C. of V.	0·657669	0·722295 (0·256)	0·749084 (0·362)	0·751322 (0·371)	0·751364 (0·371)	1·163387 (2·002)	0·910267
Theil	0·173988	0·203110 (0·291)	0·206373 (0·324)	0·207253 (0·333)	0·207334 (0·334)	0·370234 (1·963)	0·273954
Atkinson							
0·5	0·083780	0·100057 (0·301)	0·100786 (0·315)	0·101042 (0·319)	0·101110 (0·321)	0·159538 (1·402)	0·137812
1·0	0·160935	0·199542 (0·303)	0·200233 (0·309)	0·200564 (0·311)	0·200776 (0·313)	0·284522 (0·971)	0·288211
3·0	0·376905	0·543007 (0·464)	0·535410 (0·443)	0·544354 (0·468)	0·545582 (0·472)	0·615138 (0·666)	0·734553

Notes:
1. The 19 income classes of Table III have been merged into 5 classes.
2. The interpolated frequency function become negative in the last two intervals for f_6 and the last three intervals for f_7; f_8 was therefore employed in these intervals.

measures. Although the lower and upper bounds have now moved much wider apart, the interpolated inequality measures remain quite close to those of Table III. The D-measure also holds up to expectations, being still approximately $\frac{2}{3}$ for the Gini and $\frac{1}{3}$ for the other inequality measures.[21] Thus these various interpolation methods do produce robust inequality measures, more or less irrespective of the assumptions we make about the income distribution in the midrange.

Does this conclusion also apply to the end intervals?[22] To examine this we first increased the income range to include all twenty income classes. The inclusion of the lowest income class increases *all* the values of the inequality measures and, of course, widens the bounds. Nevertheless, the D-measure for the Gini coefficient remains close to $\frac{2}{3}$ for all interpolation methods, although the approximation of the D-measure to $\frac{1}{3}$ for the other measures is not quite as good as before. Next we truncated the first 9 income classes thus reducing the total number to 11. Again, as we would expect, all the inequality measures drop in value, but now all the D-measures conform closely to the "one third/two thirds" rule. Finally, we examined the top income class, reducing the assumed upper bound to 1,500,000 kr. and then increasing it to 4,000,000 kr. The difference observed either in the new lower and upper bounds or in the inequality measures based on the interpolation methods was negligible and the results were almost identical to those of the "Standard Case", the only exception being the coefficient of variation which, as we know is very sensitive to the dispersion of the data in the upper tail.

As we have seen, the inequality measure estimated either from the "bounds" or from the interpolation methods is sensitive to the particular assumption made about the lower tail of the income distribution. However, with this one exception, our results have clearly shown that the various methods of interpolation produce inequality measures which are robust with regard to a wide range of assumptions about the income distribution. The closeness of the values obtained from these various methods suggests that there is really not much to choose between them and that one might just as well stick to the simplest one, the "split histogram" method. In this regard, it is interesting to compare method 8 with method 9 which we have neglected so far—"split" and "simple" histograms. Simple histograms, to which students of applied statistics are introduced at a very early stage, produce awful results—see for example Table III—because the method is inconsistent with the condition that the mean of each interval should be equal to μ_θ. Split histograms on the other hand only need a tiny amount more work and produce amazingly good results.

However the D-measures obtained in our results show that one can go even further than this and simply estimate the lower and upper bounds for each inequality measure, then use the $\frac{2}{3}$ rule to get an estimate for the Gini measure, and use the $\frac{1}{3}$ rule for the others.

7. SUMMARY AND CONCLUSIONS

We have demonstrated two remarkable results that should eliminate a lot of the pother concerning the estimation of inequality measures from grouped data.

1. Interpolation in empirical income distributions can be performed just as well by a simple discontinuous function as by more elaborate methods. The simplest, universally feasible, such function is the "split histogram"—the step function illustrated in Figure 2 which provides as good results as more sophisticated functions but with *much* less complications in computation. All one has to do within any interval θ is to let the frequency (histogram height) be $k_\theta l_\theta$ to the *left* of the group mean and k_θ / l_θ to the *right* of the group mean where $k_\theta \equiv n_\theta / n[a_{\theta+1} - a_\theta]$ and $l_\theta \equiv [a_{\theta+1} - \mu_\theta]/[\mu_\theta - a_\theta]$.

2. In most cases the "one third/two thirds" rate produces superb results anyway, and one may disperse with even the minimal computation required for split histograms. For the Gini coefficient this means one takes $\frac{2}{3}$ times the upper bound plus $\frac{1}{3}$ times the lower bound; for *every* other NODI measure (which means virtually every other inequality measure worth considering in this context) one takes $\frac{1}{3}$ times the upper bound plus $\frac{2}{3}$ times the lower bound.

Of course there is a qualification to be made to these strong results since there are departures from the "one third" rule under some assumptions concerning the lowest income interval. Here the problem is that the upper grouping bound is obviously remarkably sensitive to the choice of a_1 when inequality measures with high inequality aversion are used. This illustrates the care which must be taken over any arbitrary assumption about the lower limit of this interval rather than any basic weakness in our simple interpolation rules. Indeed it is interesting to note that under *all* the experiments carried out the split histogram (method 8) performed as well as any of the other methods of interpolation. By contrast, although the particular assumption about the assumed upper bound to the top interval can significantly affect the upper bound of the inequality measure, nevertheless the interpolated inequality measures (with the possible exception of the coefficient of variation) remain remarkably robust. In virtually all cases the "two thirds" rule for the Gini remains unchallenged.

There are three important lessons for researchers into inequality, further discussion of which is beyond the scope of this paper. Firstly if one chooses measures with high inequality aversion, this may entail unreliable and possibly unusable empirical estimates from grouped data. Secondly, as we illustrated in Section 5, where the data are from a sample survey, for some inequality measures the sampling error may be so large as to swamp grouping error, and may render *any* interpolation method ineffective. Thirdly, without some consistent and feasible method of interpolation, empirical inequality trends may be meaningless where the interval grid changes (look at Tables III and IV).[23]

Finally, there is an important lesson for compilers of official statistics. This isn't the usual plea for more data (however welcome that may be). Presumably for most statistical bureaux the cost of publishing data is an increasing function of the number of numbers. What we have seen is that a very reliable result can be found using relatively few published numbers if they are the right ones. So if one has to make the choice between producing detailed information about either the middle income ranges or about each tail, the latter should be chosen. Moreover if the choice is between producing a breakdown of the distribution by 20 intervals without the interval means or a breakdown by 10 intervals with the interval means, again the latter option should be taken.

APPENDIX A

Properties of the 'split histogram' method

It has been noted in the text that in principle the histogram may be split at any arbitrary point in $[a_\theta, a_{\theta+1})$. For convenience we shall take the point of split as $b_\theta = \mu_\theta$

$$\begin{aligned} f_8(y) &= \frac{n_\theta}{n}\frac{a_{\theta+1}-\mu_\theta}{[a_{\theta+1}-a_\theta][\mu_\theta-a_\theta]} \equiv p, \qquad y\in[a_\theta,\mu_\theta) \\ &= \frac{n_\theta}{n}\frac{\mu_\theta-a_\theta}{[a_{\theta+1}-a_\theta][a_{\theta+1}-\mu_\theta]} \equiv q, \qquad y\in[\mu_\theta, a_{\theta+1}). \end{aligned} \tag{A.1}$$

For convenience write $a_\theta \equiv a$, $a_{\theta+1}\equiv b$, $\mu_\theta = m$, let the true density be $\psi(y)$, and consider the class of NODI measures that is ordinally equivalent to $\int_0^\infty h(y)f(y)\,dy$ where $h(\cdot)$ is a convex function.

Consider the departure of $f_8(y)$ from $\psi(y)$: we can express the resulting D-measure as a linear combination of the departure of $f_8(y)$ from $\psi(y)$ in each interval $\theta, = 1, 2, \ldots, \omega$ so we may take one such interval and look at its contribution to the D-measure resulting from the use of $f_8(y)$ rather than $\psi(y)$. That quantity may be written

$$\delta = \int_a^b h(y)[f_8(y)-\psi(y)]dy \tag{A.2}$$

$$= \int_0^{b-m} h(y+m)[q-\psi(y+m)]dy - \int_0^{a-m} h(y+m)[p-\psi(y+m)]dy \tag{A.3}$$

using (A.1). Evaluating (A.3) using a Taylor expansion we have

$$\begin{aligned} \delta &= [b-m]h(m)[q-\psi(m)]-[a-m]h(m)[p-\psi(m)] \\ &\quad + \frac{[b-m]^2}{2!}[h'(m)[q-\psi(m)]-h(m)\psi'(m)] \\ &\quad + \frac{[a-m]^2}{2!}[h'(m)[p-\psi(m)]-h(m)\psi'(m)]+[\cdots]. \end{aligned} \tag{A.4}$$

Using the definitions of p and q in (A.1), we find that (A.4) yields

$$\delta = h(m)\left[\frac{n_\theta}{n} - \psi(m)[b-a]\right] + [b-a]\left[m - \frac{a+b}{2}\right][h(m)\psi'(m) + \psi(m)h'(m)] + [\cdots]. \tag{A.5}$$

In order to examine the expansion (A.5), consider first of all the neglected higher-order terms. If we restrict attention to scale-independent NODI measures, $h(y) = y^{\beta+1}/\beta[\beta+1]$, and a, b, m may be written as proportions of the mean μ. As long as h'', ψ'' remain bounded we may neglect the higher order terms. For $0 < a < b < 1$ this may present no problem, except for some frequency distributions that peak sharply somewhere below the mean. For $a > 1$, there may be a problem with convergency in the top interval if it is open. However, even here it may be legitimate to neglect the higher order terms. For example if $\psi(y)$ is Paretian, $h(y)\psi(y)$ will be proportional to $y^{\beta-\alpha}$, and it can be seen that convergence of higher order terms will occur if $\beta < -2$.

Leaving aside the problem of the top, open interval (a problem which, as we have seen occurs with many interpolation methods), let us examine the likely sign and magnitude of δ. Clearly δ will be small

(i) if the interval width is small,
(ii) if the interval mean is located near the midpoint,
(iii) if $d/dy\ (h(y)\psi(y))$, evaluated at the interval mean, is small,
(iv) if the frequency at the interval mean approximately equals the height of the regular histogram.

Now if $\psi'(y) < 0$ throughout the interval, $m < [a+b]/2$ and hence $\psi(m) > \psi([a+b]/2$. Furthermore, if $\psi(y)$ is concave over the interval, $\psi([a+b]/2) > (n_\theta/n)/[b-a]$. Hence if $\psi' < 0$ *and* $\psi'' < 0$, $\delta < 0$. Conversely, if $\psi' > 0$ *and* $\psi'' > 0$, $\delta > 0$. So if $\psi(m)h'(m) + h(m)\psi'(m)$ is positive. The underlying distribution has a smooth, unimodal shape, the split histogram will underestimate inequality in intervals just to the right of the mode, and will overestimate inequality in the bottom intervals. Otherwise the combined effect will be ambiguous.

The D-measure for NODI measures under the split histogram

We begin with additive functions of the form (14), which we know can be decomposed as (16). Consider inequality within an arbitrary interval $[a, b)$ for which the mean is m. We know from (2) that maximum inequality within the interval is

$$I_2^* = \frac{1}{\beta+\beta^2}\left[\frac{b-m}{b-a}\left[\frac{a}{m}\right]^{\beta+1} + \frac{m-a}{b-a}\left[\frac{b}{m}\right]^{\beta+1} - 1\right]. \tag{A.6}$$

Since we are taking a single interval, mininum inequality is evidently zero. If we assume that distribution 8—the split histogram—holds then integration over $[a, \mu)$ and $[\mu, b)$ reveals that inequality within the interval is

$$I_8^* = \frac{1}{\beta+\beta^2}\left[\frac{m}{[\beta+2][b-a]}\right]\left[\frac{b-m}{m-a}\left[1-\left[\frac{a}{m}\right]^{\beta+2}\right] + \frac{m-a}{b-m}\left[\left[\frac{b}{m}\right]^{\beta+2} - 1\right] - 1\right]. \tag{A.7}$$

Introduce the variables $x \equiv 1-(a/m)$, $z \equiv (b/m)-1$; clearly $0 < x < 1$ and $0 < z$ for non-trivial distributions. We shall also assume that the interval width is not too large, so that $b < 2m$. Substitution in (A.6) and (A.7) and a little rearrangement yields, respectively

$$I_2^*[\beta+\beta^2][x+z] = z[[1-x]^{\beta+1} - 1] + x[[1+z]^{\beta+1} - 1] \tag{A.8}$$

$$I_8^*[\beta+\beta^2][x+z]=z\left[\frac{1-[1-x]^{\beta+2}}{x[\beta+2]}-1\right]+x\left[\frac{[1+z]^{\beta+2}-1}{z[\beta+2]}-1\right]. \tag{A.9}$$

Expanding (A.8) yields

$$z\left[1-[\beta+1]x+\beta[\beta+1]\frac{x^2}{2!}-[\beta-1]\beta[\beta+1]\frac{x^3}{3!}+\cdots-1\right]$$
$$+x\left[1+[\beta+1]z+\beta[\beta+1]\frac{z^2}{2!}+[\beta+1]\beta[\beta-1]\frac{z^3}{3!}+\cdots-1\right]. \tag{A.10}$$

Simplifying (A.10) we see that (A.8) can be rewritten

$$I_2^*=zx+[\beta-1]\frac{zx[z-x]}{3!}+\cdots. \tag{A.11}$$

Expanding (A.9) likewise yields

$$z\left[1-[\beta+1]\frac{x}{2!}+\beta[\beta+1]\frac{x^2}{3!}-[\beta-1]\beta[\beta+1]\frac{x^3}{4!}+\cdots-1\right]$$
$$+x\left[1+[\beta+1]\frac{z}{2!}+\beta[\beta+1]\frac{z^2}{3!}+[\beta-1]\beta[\beta+1]\frac{z^3}{4!}+\cdots-1\right]. \tag{A.12}$$

Simplifying (A.12) we see that (A.9) can be rewritten

$$I_8^*=\frac{zx}{3}+[\beta-1]zx\frac{[z-x]}{4!}+\cdots. \tag{A.13}$$

Clearly for small z and x and for β close to unity the last term in (A.11) and in (A.13) becomes negligible. In particular, if the distribution within the interval has its mean at the midpoint $(a+b)/2$, then $z=x$ and all the even terms in the series (A.10) and (A.12) vanish. Under such circumstances we see immediately that $I_8^*\simeq\frac{1}{3}I_2^*$. Hence if *every* interval has the split histogram distribution the D-measure given in (17) must be approximately $\frac{1}{3}$.

Now consider the Gini coefficient which is decomposable as in (15). Evaluating maximum inequality within the interval $[a, b]$ using (A.1) we find

$$I_2^*=\frac{[m-a][b-m]}{m[b-a]}. \tag{A.14}$$

Now consider Gini inequality in $[m, b)$ and the Gini that would result if all the population in $[a, b)$ were to be concentrated in the appropriate proportions at $\frac{1}{2}[a+m]$ and $\frac{1}{2}[m+b]$. These correspond to the three terms on the RHS of the expression below

$$I_8^*=\left[\frac{b-m}{b-a}\right]^2\frac{[m-a]}{m}+\left[\frac{m-a}{b-a}\right]^2\frac{b-m}{m}+\frac{[b-m][m-a]}{2m[b-a]}. \tag{A.15}$$

Evaluating (A.15) we get immediately

$$I_8^*=\frac{3[b-m][m-a]}{2m[b-a]}$$

whence $I_8^*=\frac{2}{3}I_2^*$. Hence if every interval has the split histogram distribution, the D-measure given in (17) must be *exactly* $\frac{2}{3}$.

Estimation procedures

To compute A_θ, α_θ in (5) note that in order to satisfy conditions (a) and (b) we must have

$$\frac{n_\theta}{n}=A[a_\theta^{-\alpha}-a_{\theta+1}^{-\alpha}] \tag{A.16}$$

$$\mu_\theta=\frac{\alpha}{\alpha-1}[a_\theta^{1-\alpha}-a_{\theta+1}^{1-\alpha}] \tag{A.17}$$

the parameter values are then found as the values of A, α which satisfy (A.16) and (A.17).

To calculate $\gamma_{\theta 0},\ldots,\gamma_{\theta 4}$ in (6) the following conditions on the spline were imposed. (i) Conditions (a) and (b) must be satisfied in interval θ—this yields two equations like (A.16) and (A.17) for the quartic interpolation. (ii) The right-hand ordinate at $a_{\theta+1}$ must equal the previously computed left-hand ordinate from interval $\theta+1$, the spline being commenced from the upper tail with knots preassigned at the interval boundaries. Where the assumption of a finite $a_{\omega+1}$ was used it was assumed that $f_6(a_{\omega+1})=0$. (iii) The right-hand derivative of f_6 at $a_{\theta+1}$ must equal the previously computed left-hand derivative from interval $\theta+1$. If the top interval is finite we assume $f_6'(a_{\omega+1})=0$. Conditions (ii) and (iii) determine two further equations. (iv) In order to reduce the chance of the spline intersecting the horizontal axis in interval θ, the extension of the spline function into interval $\theta-1$ was constrained to go through the point $([a_{\theta-1}+a_\theta]/2,\ n_{\theta-1}/n[a_\theta-a_{\theta-1}])$. It was found that cubic splines, which of course do not use this fourth restriction, frequently violated the horizontal axis and were therefore unsatisfactory. A check was routinely made for violation of the horizontal axis by the quartic spline; on the rare occasions where this occurred, linear interpolation was used instead.

Linear estimates were found by solving for $\gamma_{\theta 0}$ and $\gamma_{\theta 1}$ from the equations for n_θ/n and μ_θ analogous to (A.16), (A.17). If the straight line intersected the horizontal axis, the split histogram was used.

TABLE V

Income groups ('000 kroner)	Total income for each interval	Total tax payers in the interval
0	–	118,073
0·1–4·9	396,477	155,842
5·0–9·9	2,081,416	278,416
10·0–14·9	6,880,865	539,659
15·0–19·9	9,471,991	544,600
20·0–24·9	9,557,162	425,362
25·0–29·9	11,065,055	402,923
30·0–34·9	12,420,679	382,345
35·0–39·9	13,801,434	368,096
40·0–44·9	17,025,900	400,028
45·0–49·9	21,706,656	456,660
50·0–54·9	24,163,123	460,530
55·0–59·9	22,013,271	383,588
60·0–69·9	30,933,655	480,228
70·0–79·9	17,912,913	240,578
80·0–99·9	18,529,954	209,900
100·0–119·9	9,018,675	83,030
120·0–149·9	6,618,096	49,969
150·0–199·9	4,792,727	28,149
200·0–499·9	4,184,993	16,287
500·0	773,818	892

APPENDIX B

The data used in Section 3 of this paper are presented in Table V. The variable we have used is total income for the Income Year 1977. Column 3 indicates frequency of each income class. Column 2 is the sum of the incomes of those falling in each particular income class. The source of these data is Table 1 (p. 22) of Statistiska Meddelanden (N1978:22) published by the National Central Bureau of Statistics, Fack, S102 50 Stockholm, Sweden.

First version received August 1980; *final version accepted October* 1981 (*Eds.*).

We are grateful to the Suntory–Toyota foundation and the SSRC for financial support.

NOTES

1. See for example Cowell (1977), Gastwirth (1972, 1975).
2. Surprisingly this issue is often ignored. The non-negative frequency requirement is formally equivalent to requiring that the Lorenz curve be convex. For example, Kakwani's (1976) empirical work does not deal with this fundamental issue. See also Gastwirth and Glauberman (1976) who do take account of this problem.
3. See Gastwirth (1972, 1975).
4. See for example Gastwirth and Glauberman (1976), Needleman (1978), Kakwani (1976).
5. The most convenient assumption is that $b_\theta = \mu_\theta$, since then Property 1 always holds. We have made this assumption throughout.
6. See Appendix A.
7. In Section 6 we examine a number of possible values for L including two extreme cases. See also Needleman (1978).
8. All the calculations reported in Section 6 were also performed using the assumption of a Paretian upper tail. However, since the results were so similar to these in Tables I, III, and IV we have not reported them separately.
9. See Cowell and Shorrocks (1980).
10. See Cowell (1980), Cowell and Kuga (1981*a*, *b*) which also discuss the special cases $\beta = 0$, $\beta = -1$.
11. This is similar to the goodness-of-fit test suggested by Gastwirth and Smith (1972).
12. In Appendix A we examine this rigorously for f_8, the split histogram distribution. We are indebted to Professor D. G. Champernowne who first suggested the $\frac{2}{3}$ and $\frac{1}{3}$ rules in an unpublished manuscript, and who has shown that the $\frac{1}{3}$ rule will be approximately true for a number of measures of the I^β type in the case of a "cubic" interpolated Lorenz curve.
13. It is interesting to note, however, that in the large sample results reported by Gastwirth (1972), the "$\frac{2}{3}$ rule" for the Gini is confirmed.
14. For further confirmation see *A Panel Study of Income Dynamics (1978), Procedure and Tape Codes 1977 Interviewing Year, WAVE X, A Supplement,* Institute for Social Research, University of Michigan.
15. The standard errors for the inequality measures I^β were derived from standard formulae on the standard errors of moments, and the standard error for I^G was found from a conventional approximation using the assumption of asymptotic normality—see Kendall and Stuart, pp. 228–241. Other interpolation methods were also utilized, but since the results were so similar to those for f_5 and f_8, these are not separately reported. Moreover, we computed the Atkinson index for a wide range.
16. Other "non-NODI" measures such as the logarithmic variable and the relative mean deviation were also computed but are not reported here. Similar results are obtained, however. For further discussion of all of these measures see Cowell (1977).
17. We had tried other groupings as well, but the results are unchanged and therefore not reported.
18. To see this in perspective, note that the sample size of the U.S. Current Population Survey is "only" of the order of 50,000!
19. See Appendix B for a description of the data source.
20. Note that in Tables III and IV the first and last columns are the lower and upper bounds I_1, I_2 respectively.
21. The obvious exception is the Atkinson measure with inequality aversion parameter equal to 3. However such a value for this parameter is really very high and implies great sensitivity to the detail of what happens within the lower interval. We have included this as an extreme case of the NODI measures, and it is to be expected that here, and in the analysis of the sensitivity of the bottom interval, significant departures from the $\frac{1}{3}$ rule are evident for inequality aversion values of 3 and above. For values less than 3, the associated D-measures conformed to those for other inequality measures.
22. Tables of the results referred to in this paragraph are available from the authors on request.
23. This issue has also been examined by Petersen (1979).

REFERENCES

COWELL, F. A. (1977) *Measuring Inequality* (Oxford: Philip Allan Publishers Ltd.).

COWELL, F. A. and KUGA, K. (1981*a*), "Inequality Measurement—An Axiomatic Approach", *European Economic Review* **15**, 287–305.

COWELL, F. A. and KUGA, K. (1981*b*), "Additivity and the Entropy Concept: An Axiomatic Approach to Inequality Measurement", *Journal of Economic Theory*, **25**, 131–143.

COWELL, F. A. and SHORROCKS, A. F. (1980), "Inequality Decomposition by Population Subgroups" (LSE mimeo).

GASTWIRTH, J. L. (1972), "The Estimation of the Lorenz Curve and Gini Index", *Rev. Econ. and Statis.*, **63** (3), 306–316.

GASTWIRTH, J. L. (1975), "The Estimation of a Family of Measures of Economic Inequality", *Journal of Econometrics*, 61–70.

GASTWIRTH, J. L. and GLAUBERMAN, M. (1976), "The Interpolation of the Lorenz Curve and Gini Index from Grouped data", *Econometrica*, **44** (3), 479–483.

GASTWIRTH, J. L. and SMITH, J. T. (1972), "A new Goodness-of-fit Test", *Business & Economic Section Proceedings of the American Statistical Association*, 320–322.

KAKWANI, N. (1976), "On the Estimation of Income Inequality Measures from Grouped Observations", *Review of Economic Studies*, **43**, 483–492.

KENDALL, M. G. and STUART, A. (1958) *Advanced Theory of Statistics* (London: Charles Griffith & Co. Ltd.) Vol. 1.

NEEDLEMAN, L. (1978), "On the Approximation of the Gini Coefficient of Concentration", *The Manchester School*, 105–122.

PETERSEN, H.-G. (1979), "Effects of Growing Incomes in Classified Income Distributions, the Derived Lorenz Curves and Gini Indices", *Econometrica*, **47**, 183–198.