Inequality: Measurement

Frank A. Cowell
STICERD
London School of Economics

DARP 86
August 2006

The Toyota Centre
Suntory and Toyota International
Centres for Economics and Related
Disciplines
London School of Economics
Houghton Street
London WC2A 2A

(+44 020) 7955 6674

Abstract

This article provides a brief overview of the key issues in inequality measurement and has been prepared for inclusion in the second edition of The New Palgrave.

My thanks go to Yoram Amiel, Tony Atkinson, Sanghamitra Bandyopadhyay, Kristof Bosmans, Udo Ebert, Giovanni Ko, Peter Lambert and Abigail McKnight who made helpful comments on an earlier draft.

JEL Classification: C13, D63

Keywords: inequality, social welfare, ranking.

Distributional Analysis Research Programme

The Distributional Analysis Research Programme was established in 1993 with funding from the Economic and Social Research Council. It is located within the Suntory and Toyota International Centres for Economics and Related Disciplines (STICERD) at the London School of Economics and Political Science. The programme is directed by Frank Cowell. The Discussion Paper series is available free of charge. To subscribe to the DARP paper series, or for further information on the work of the Programme, please contact our Research Secretary, Leila Alberici on:

Telephone: UK+20 7955 6674
Fax: UK+20 7955 6951
Email: l.alberici@lse.ac.uk
Web site: http://sticerd.lse.ac.uk/DARP

1 Introduction

Inequality measurement is principally concerned with the comparison of personal income distributions in quantitative terms. In its modern form it is a branch of welfare economics although it clearly derives some of its intellectual heritage from statistics. It is distinct from the measurement of poverty and relative deprivation although there are close analytical links to these topics. The motivation for taking the subject of inequality seriously is both analytical and practical: the principal concepts reviewed in this article are of concern to theoretical economists and are also used by policy makers. The subject touches on questions addressed by philosophers and by social scientists.

q	q-quantile		Growth
	1974	**2004**	
10%	\$9,741	\$10,927	12.2%
20%	\$16,285	\$18,500	13.6%
50%	\$37,519	\$44,389	18.3%
80%	\$64,781	\$88,029	35.9%
90%	\$83,532	\$120,924	44.8%
95%	\$102,534	\$157,185	53.3%

Note: Cols 2, 3 give the upper limit of the bottom 10%, 20%,... of the population

Table 1: Quantile incomes and growth. US 1974-2004

The type of issue under consideration can be illustrated by a simple example as depicted in Tables 1 and 2. These tables do not pretend to be the most general or the most suitable representation of the facts, but they are from an easily accessible source[1] and give a convenient snapshot of what happened to the distribution of income in the United States over a span of about thirty years. From Table 1 it is clear that the bottom decile income experienced a 12.2 percent growth over the period (in real terms) while the median grew by half as much again (18.3%) and the top decile grew by almost four times as much (44.8%). Table 2 describes what happened to the average incomes of particular *groups*. The average income of households in the bottom fifth of the distribution grew by just 10.1% over the thirty years while the average income of households in the top fifth grew by 58.6%. More on using the concepts of quantiles and shares once we have introduced some of the technical equipment needed for analysing income distributions.

The thumbnail sketch suggests a substantial increase in inequality in the United States over the last quarter of the 20th century. But how much did inequality increase? In what ways can the impressionistic method of inequality comparisons suggested in the example be made precise and interpreted within the context of standard economic analysis? The purpose of this article is to

[1] Data are from DeNavas-Walt, Proctor, and Lee (2005) Appendix Table A3 available at http://www.census.gov/prod/2005pubs/p60-229.pdf. Incomes are in 2004 dollars; the income-receiving unit is the household. See page 29 for details on the way income data are assembled.

Group	Average income		Growth
	1974	**2004**	
1st	\$9,324	\$10,264	10.1%
2nd	\$23,176	\$26,241	13.2%
3rd	\$37,353	\$44,455	19.0%
4th	\$53,944	\$70,085	29.9%
top	\$95,576	\$151,593	58.6%
Overall	*\$43,875*	*\$60,528*	*38.0%*

Note: Cols 2, 3 give the average incomes of the bottom fifth, second fifth, ...

Table 2: Growth in average incomes for the five quintile groups and overall. US 1974-2004

provide a succinct overview of the role played by economic theory and other abstract principles in this class of problem and how to make sense of inequality comparisons such as those suggested in the example.

The sketch example in Tables 1 and 2 also illustrates some of the essential practicalities that have to be taken into account when implementing the principles of inequality measurement: should we be focusing on households or individuals? What is the appropriate definition of income?

To follow the analysis there are few prerequisites: an understanding of utility and preference analysis is helpful but not essential to grasping the basic points that will be discussed.

2 Basics

2.1 Components of the problem

The framework adopted here is not the most general approach, but one that is suitable for setting out the key ideas. We begin by considering the basic building blocks and then show how to assemble the constituent parts.

2.1.1 Income and income distribution

At the heart of the problem there is some scalar entity to be called "income," but in practice this entity could be wealth, expenditure or some other economic quantity, the distribution of which is of particular interest. Income is distributed among a number of "income receivers," which we will refer to as "persons" (although the income receiver in practice may be a family or household). Suppose that there is a known number of income receivers n and that person i has income x_i. The *income distribution* is then simply the vector

$$\mathbf{x} = (x_1, x_2, ..., x_n). \tag{1}$$

The set of all possible income distributions X is a subset of $\mathbb{R}^n$. The nature of X is going to depend in practice upon the precise definition of "income:" is it logically possible to have a zero value of x_i, for example? Or a negative value? As a working assumption we will take it that X consists of all vectors (1) such that $x_i \geq \underline{x}$ and leave open the specification of the lower bound $\underline{x}$ for particular instances of the inequality-measurement problem. Representations of the income distribution other than (1) will appear later in the discussion.

2.1.2 Indices

The topic of "inequality measurement" presumes that there is an inequality measure. An obvious interpretation of this is that there is some index I that, given a particular income distribution $\mathbf{x}$, yields a real number that is taken to be the amount of inequality exhibited by the distribution. In some ways the index I works like other well-known summary statistics of distributions, such as the mean

$$\mu(\mathbf{x}) := \frac{1}{n}\sum_{i=1}^{n} x_i \tag{2}$$

and the variance

$$\text{var}(\mathbf{x}) := \frac{1}{n}\sum_{i=1}^{n} \left[x_i - \mu(\mathbf{x})\right]^2. \tag{3}$$

Indeed the variance itself is sometimes used as an inequality index, although it is more common to use a transformed version of it known as the *coefficient of variation*:

$$I_{\text{CV}}(\mathbf{x}) := \frac{\sqrt{\text{var}(\mathbf{x})}}{\mu(\mathbf{x})}. \tag{4}$$

One of the most commonly used indices in practice is the *Gini coefficient* defined as

$$I_{\text{Gini}}(\mathbf{x}) := \frac{1}{2n^2\mu(\mathbf{x})}\sum_{i=1}^{n}\sum_{j=1}^{n} \left|x_i - x_j\right|. \tag{5}$$

There are many more. However, rather than running through an exhaustive list of candidate indices it is more useful to examine the principles that have usually been applied to construct indices; this we do by considering *a priori* what constitutes a "suitable" inequality measure, the issue addressed in section 2.2.

2.1.3 Ranking and dominance

An apparently more flexible interpretation of the idea of inequality measurement is the idea of an inequality *ranking*. This is a partial ordering that picks up the general flavour of the kind of comparisons that we suggested in the introduction; the partial ordering is typically captured by a simple representational tool. Consider three of these.

The first of these tools is *Pen's parade* (named after the famous parable introduced by Pen (1974)), which is simply the inverse of the empirical distribution function. To depict it let $x_{[i]}$ denote the ith smallest component in the vector (1) – the ith smallest income. Then take the collection of points

$$\left(\frac{i}{n}, x_{[i]}\right), \; i = 1, 2, ..., n. \tag{6}$$

From this simple definition we can also introduce the the idea of dominance. Take two distributions $\mathbf{x}'$ and $\mathbf{x}''$ in X where $\mathbf{x}' = (x'_1, x'_2, ..., x'_n)$ and $\mathbf{x}'' = (x''_1, x''_2, ..., x''_n)$. If it is true that $x'_{[i]} > x''_{[i]}$ for all $i = 1, 2, ..., n$ then we say that $\mathbf{x}'$ strictly Parade-dominates $\mathbf{x}''$.

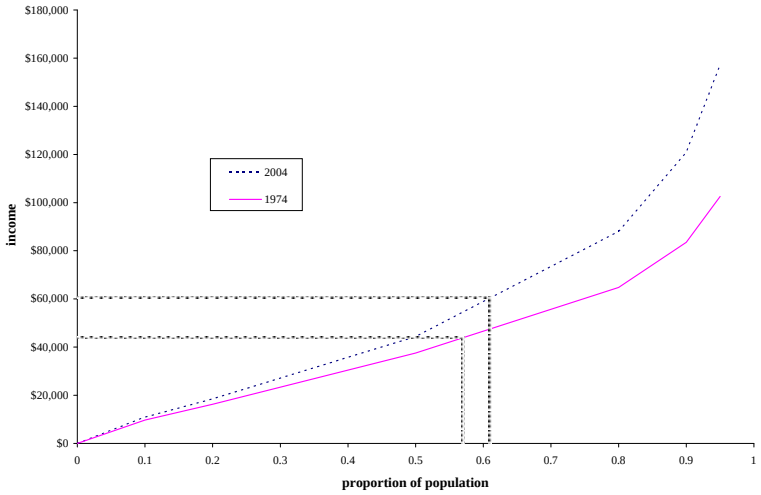

Figure 1: Parade diagram corresponding to Table 1.

The resulting graph plots income quantiles against population proportions: $x_{[i]}$ is the quantile corresponding to the bottom q percent of the population where $q = 100\frac{i}{n}$. To illustrate the concept we use the information in Table 1 to produce a graph that looks like Figure 1. In Pen's parable we imagine the whole population (seen as individuals rather than households) arranged in order on the $[0, 1]$ interval where each person's height has been altered in proportion to his/her income; the average-height income recipient in 1974 is located at position 0.57 in Figure 1[2] but in 2004 the average-height income recipient is located at position 0.61. Although the distribution of 2004 Parade-dominates the distribution in 1974, it is clear from Table 1 that overall the Parade shifted upwards in a lopsided fashion over the thirty years with the

[2] In other words at a point 57% along the horizontal axis the height of the Parade is exactly mean income.

incomes of the very rich (95% quantile) growing more than four times faster than the poor (10% quantile); this shift suggests increased inequality over the period. However, by itself the Parade does not tell us much about inequality directly, although concepts closely related to it are widely used to characterise inequality comparisons. It is common to use *quantile ratios* for distributional comparisons: for example the popular "90-10 ratio" is given by $x_{[k]}/x_{[j]}$ where j and k are, respectively, the smallest integers satisfying $j/n \geq 10\%$ and $k/n \geq 90\%$: in the example above this ratio increased from 8.6 to 11.1. Furthermore there is an important welfare-economic interpretation of the Parade that is discussed in section 3.3 below.

For the second and third concepts we use the $x_{[i]}$ to derive the normalised income cumulations; for any $i = 1, 2, ..., n$ these are

$$c_i := \frac{1}{n} \sum_{j=1}^{i} x_{[j]}. \tag{7}$$

Then the *generalised Lorenz curve* (GLC) is given by the graph of

$$\left(\frac{i}{n}, c_i \right), \ i = 1, 2, ..., n. \tag{8}$$

Again we have a natural definition of dominance: for two distributions $\mathbf{x}'$ and $\mathbf{x}''$ in X if it is true that $c'_i > c''_i$ for all $i = 1, 2, ..., n$ then we say that $\mathbf{x}'$ strictly GLC-dominates $\mathbf{x}''$.[3] For the example we used earlier the GLC is illustrated in Figure 2, derived from Table 2.

The GLC plots the normalised income of the bottom $100q$ percent of the population against q and, although the 2004 distribution GLC-dominates 1974, it is clear that over the period the growth of these group averages was not evenly distributed – the higher was q the higher was the growth over 1974 to 2004.[4] Once again, although the GLC does not give information about inequality comparisons directly, there is an important welfare-economic interpretation (in section 3.3). In addition a small modification of the GLC yields one of the central concepts of distributional analysis. Dividing c_i in (7) by the mean $\mu(\mathbf{x})$ gives the income share of the bottom $100\frac{i}{n}$ percent of the population. The graph of the (population-proportion, income-share) pairs

$$\left(\frac{i}{n}, \frac{c_i}{\mu(\mathbf{x})} \right), \ i = 1, 2, ..., n \tag{9}$$

gives the *Lorenz curve*. Also, for two distributions $\mathbf{x}'$ and $\mathbf{x}''$, if it is true that, $c'_i/\mu(\mathbf{x}') > c''_i/\mu(\mathbf{x}'')$ for all $i = 1, 2, ..., n-1$ then we say that $\mathbf{x}'$ strictly Lorenz-dominates $\mathbf{x}''$. In the case of the example using US data this is illustrated in

[3] Note that the definitions of Parade- and GLC-dominance can be extended to cases where the two distributions do not have the same number of incomes – this step makes use of the "population principle" defined in Section 2.2 below. In some cases it is useful to consider the weak (non-strict) versions of the dominance criteria introduced here.

[4] This is easily inferred from Table 2: for example the average income of the top 20% grew almost six times as fast as the average income of the bottom 20%.

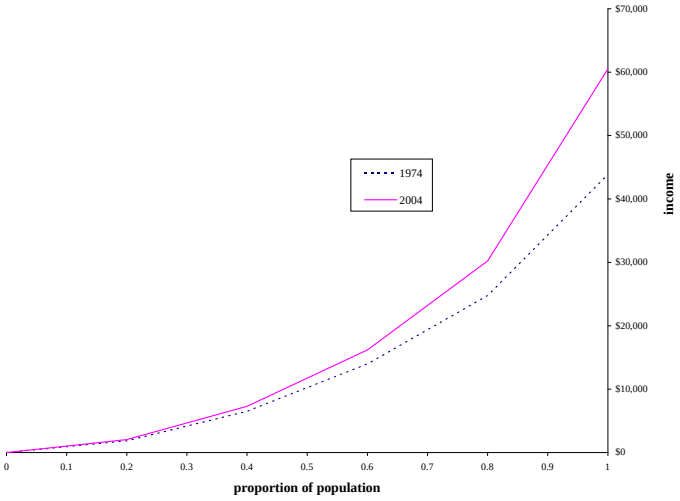

Figure 2: Generalised Lorenz curve: Source as for Table 1

Figure 3: the Lorenz curve plots the income share of the bottom $100q$ percent of the population against q and the diagonal line depicts a hypothetical distribution of perfect equality.[5] It is clear that for each q the share was smaller in 2004 than it was in 1974 – the 1974 distribution Lorenz-dominates that for 2004. This simple intuitive notion of greater inequality conforms exactly with a fundamental principle to be explained below.

2.2 Axioms

An inequality index I is in some ways like a utility function in consumer theory: it is a representation of an inequality ordering on the members of X and is usually taken to be continuous and ordinal – although there is often a "natural" cardinal representation of a particular index, a formal argument for one representation rather than another is not usually provided (why not use the square or the log of the Gini coefficient?). Ordinality is sufficient for making comparing income distributions, the primary task of inerquality analysis. Axioms are essentially formal statements of the principles of assessment that are used to give meaning to the ordering represented by I. The treatment here does not claim to generality but rather it focuses on those principles that are central to modern approaches to inequality. Rather than presenting the axioms as formal statements, however, it is more useful here to introduce the underlying key

[5] Take the area trapped between the Lorenz curve and the equality diagonal. Using (7) and (9) we can show that the ratio of this area to the area of the whole triangle is given by the weighted sum $\sum_{i=1}^{n} \kappa_i x_{[i]}$ where the weights are $\kappa_i := [2i - 1 - n] / [n\mu(\mathbf{x})]$. This is exactly the Gini coefficient (5).

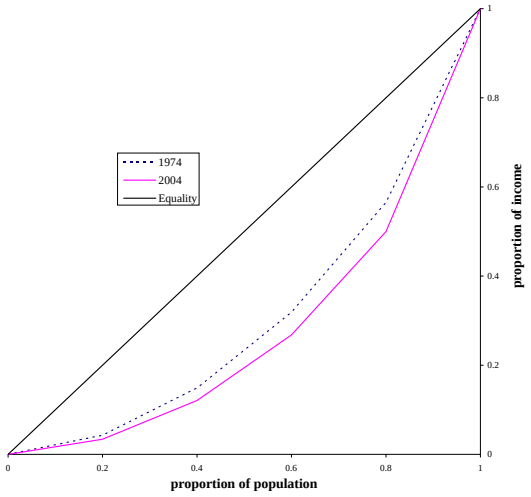

Figure 3: Lorenz Curve: Source as for Table 1

principles discursively.

Assume that everywhere in the following discussion the vector $\mathbf{x}$ in (1) is any arbitrary member of the set X.

- First, it seems reasonable that the labelling of the components of $\mathbf{x}$ be irrelevant: it does not matter which income receiver gets which income. This means that I has the *symmetry* property:

$$I(x_1, x_2, ..., x_n) = I(x_2, x_1, ..., x_n) = I(x_3, x_1, ..., x_n) = ... \qquad (10)$$

  We will always assume that this holds and we may therefore adopt the convention that incomes have been labelled such that $x_1 \leq x_2 \leq ... \leq x_{n-1} \leq x_n$.

- Second, we need some coherent way of characterising inequality in different-sized populations. Perhaps the most obvious assumption is that simple replications of an income vector (1) leave inequality unchanged. This is the *population principle*:

$$\begin{aligned} I(x_1, x_2, ..., x_n) &= I(x_1, x_1, x_2, x_2, ..., x_n, x_n) \\ &= I(x_1, x_1, x_1, x_2, x_2, x_2, ..., x_n, x_n, x_n) \\ &= ... \end{aligned} \qquad (11)$$

  Taken in conjunction with symmetry this allows one to represent distributions purely in terms of a distribution function.

- A key assumption that is commonly invoked focuses on the effect on inequality of a hypothetical small income transfer. Suppose $x_i < x_j$ and consider some positive number δ such that $x_i - \delta \geq \underline{x}$, then the *principle of transfers* (Dalton 1920) requires that:

$$I(x_1, , ..., x_i, ..., x_j, ..., x_n) < I(x_1, , ..., x_i - \delta, ..., x_j + \delta, ..., x_n) \qquad (12)$$

  – a poorer-to-richer income transfer will always increase inequality.

- As a counterpart to the assumption relating to different sizes of population (equation 11) it is useful to have an assumption relating to different amounts of total income. The standard assumption is that of *scale independence*. This requires that, for any scalar $\lambda > 0$:

$$I(\lambda x_1, \lambda x_2, ..., \lambda x_n) = I(x_1, x_2, ..., x_n) \qquad (13)$$

  – double all incomes or halve all incomes and inequality is left unaltered. An alternative assumption that is sometimes used is *translation independence*. Take any real number δ such that $x_1 + \delta \geq \underline{x}$; then

$$I(x_1 + \delta, x_2 + \delta, ..., x_n + \delta) = I(x_1, x_2, ..., x_n) \qquad (14)$$

  – add or subtract one dollar from every income and inequality is left unaltered.

Clearly this brief list raises some important questions. Why use these particular axioms? Some of them appear to be quite strong; for example, although scale independence seems attractive if the "incomes" x_i here are measured in dollars and we consider just dividing through by some rate of exchange so as to work with incomes in some other monetary units, it may seems less attractive if we want to consider the impact on inequality of redistribution policies at different stages of economic growth: a rearrangement of income shares that constitutes a reduction in inequality in a low-income society might not be considered as a reduction in inequality if the whole population is prosperous. Furthermore, the axioms captured by equations (10)–(13) for example are satisfied by both (4) and (5) as well as other important classes of inequality measures; on the other hand the axioms captured by equations (10)–(12) and (14) are satisfied by (3) and another rich class of inequality measures. Following on from this question, what more is required to get a specific index or well-defined family of indices that is both theoretically appropriate and practical to implement?

To answer this we need to be precise about what it means to say that one distribution is more unequal than another and the intellectual basis used for making such comparisons. The meaning of inequality can be further clarified through one of several routes: this article will analyse three of these in turn, namely, social welfare, decomposition, income differences.

3 Social welfare and inequality

The welfare-economic approach to the subject starts from the position that inequality is about "illfare" – the opposite of welfare. If we adopt this approach then the definition of inequality follows almost immediately. The idea is similar to the conventional measurement of economic waste and the basis for a simple model can be laid with only a little more theorising.

The *social-welfare function* (SWF) is a real-valued function W defined on the space of distributions X. The social welfare associated with a particular income distribution (1), given by

$$W(x_1, x_2, ..., x_n), \tag{15}$$

is to be interpreted as follows: suppose we are given a specific SWF $W(\cdot)$ and that for two separate income distributions $\mathbf{x}'$ and $\mathbf{x}''$ we have $W(\mathbf{x}') > W(\mathbf{x}'')$; then social welfare associated with the distribution $\mathbf{x}'$ is higher than the social welfare associated with the distribution $\mathbf{x}''$. In principle W is an ordinal function so that the scale of measurement of welfare levels can be subjected to arbitrary monotonic-increasing transformations.

This basic specification raises a number of important questions:

- Why express social welfare as a function of income? Income defined how?
- What particular form should W take?
- What is the relation between the functions I and W?

The answer to the first question helps to pin down the relationship between inequality measurement as conventionally practised and standard welfare economics – see section 3.1. The answers to the last two questions will determine the form of a class of inequality measures and permit us to establish some important welfare-economic results: these are addressed in sections 3.2 and 3.3.

3.1 Welfare and income

We need to rectify a point that was fudged in the discussion of the US example: how to do the trick of passing from a distribution of dollar income among households to a standard welfare analysis that is typically concerned with the levels of economic wellbeing of individuals. The standard approach is as follows. We require a method of appropriately capturing the relationship between the living standard that is attainable by an individual and the income that he/she is presumed to have access to within the household. This is conventionally done by defining a function $\nu(.)$ that has as its argument a list of non-income attributes $\mathbf{a}$ that might include household size, age and sex of household members and health status; $\nu(\mathbf{a})$ determines the number of *equivalent adults* in the household with attributes $\mathbf{a}$ such that

$$x = \frac{y}{\nu(\mathbf{a})}, \tag{16}$$

where y is nominal income and x is *equivalised income* that is taken to be comparable across different household types. Note that the equivalisation function ν is typically specified as independent of income although this simplification is not essential; of course the way in which the function ν is determined – from ethical considerations, or econometric studies – is an important issue in its own right, but one that lies outside the present discussion. The function ν transforms a distribution of dollar incomes among n households

$$\mathbf{y} = (y_1, y_2, ..., y_n) \tag{17}$$

into a distribution of equivalised incomes by households given by (1). In order to complete the welfare interpretation we need to recognise that social-welfare considerations are usually represented in terms of individuals rather than households and so, for example, households consisting of couples should receive more weight in social-welfare evaluations than households consisting of single individuals. Therefore, if the income-receiving units consist of households of differing size, we might want to represent this by introducing a corresponding set of population weights w_i for the observations, so that the distribution becomes an ordered list of pairs:

$$((w_1, x_1), (w_2, x_2), ..., (w_n, x_n)) \tag{18}$$

where w_i is the number of persons in household i divided by the number of persons in the whole population. There is little analytical complication in using (18) rather than (1) as a representation of the distribution of equivalised incomes by individuals. Typically it is just a matter of a minor redefinition of formulas for inequality measures and the like: for example the coefficient of variation (4) would now be written

$$\sqrt{\sum_{i=1}^{n} w_i \left[\frac{x_i}{\mu} - 1\right]^2} \tag{19}$$

where μ is the appropriately redefined mean $\sum_{i=1}^{n} w_i x_i$.[6]

However, having introduced this important theoretical qualification we will now neglect it – for expositional purposes it is convenient to assume that the population consists of isolated individuals that are identical in every relevant respect other than income and that income appropriately represents individual welfare. So, from here on, i indexes individuals or households and the distinction between x and y is dropped.

3.2 Social welfare and inequality measures

The idea of the SWF was introduced without discussing specific properties of the function W. Some properties must be imposed on W if we require there to

[6] More generally: all measures that can be written in the form $\Phi\left(\frac{1}{n}\sum_{i=1}^{n}\phi(x_i), \mu\right)$ just need to be rewritten in the form $\Phi\left(\sum_{i=1}^{n} w_i \phi(x_i), \mu\right)$. A similar modification applies to the Gini coefficient.

be a specific relationship between social welfare and inequality and we impose specific assumptions on the function I. However, in addition it is particularly important to be explicit about how W should respond to an increase in one or more incomes. This is the usual principle that is applied:

- Suppose we consider any income distribution $(x_1, x_2, ..., x_n)$ and some positive number δ. Then *monotonicity* requires that:

$$W(x_1, x_2, ..., x_i + \delta, ..., x_n) > W(x_1, x_2, ..., x_i, ..., x_n) \tag{20}$$

Assuming that monotonicity holds and that W is a continuous function, the SWF can itself be used to derive a family of inequality measures. There are several ways of doing this, but a standard approach is to represent social welfare using a money metric: we can always do this in view of the ordinal nature of W and the requirement that it be monotonic and continuous. The *equally-distributed equivalent* (EDE) income is a real number ξ such that for any $(x_1, x_2, ..., x_n)$ in X:[7]

$$W(\xi, \xi, ..., \xi) = W(x_1, x_2, ..., x_n). \tag{21}$$

Clearly the relationship (21) can be used to derive EDE as a function of the income distribution, $\xi(\mathbf{x})$ and the function $\xi(\cdot)$ is a valid way of representing social welfare.

Suppose we require that the principle of transfers apply to W; this by analogy with (12) means that a mean-preserving poorer-to-richer income transfer will *de*crease social welfare. Then it is always true that $\xi(\mathbf{x}) \leq \mu(\mathbf{x})$ and the normalised gap between ξ and μ provides a natural basis for an inequality index

$$1 - \frac{\xi(\mathbf{x})}{\mu(\mathbf{x})}. \tag{22}$$

It is clear that this index is bounded between zero and 1 and that if there were perfect equality then we would have $\xi(\mathbf{x}) = \mu(\mathbf{x})$ and inequality in (22) would be zero.

Furthermore, if the scale-independence property (13) is also satisfied, then EDE income takes the form of a generalised mean:

$$\xi(\mathbf{x}) = \left[\frac{1}{n}\sum_{i=1}^{n} x_i^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}}, \ \varepsilon > 0 \tag{23}$$

and (22) gives the class of *Atkinson indices*:[8]

$$I_{\mathrm{A}}^{\varepsilon}(\mathbf{x}) := 1 - \left[\frac{1}{n}\sum_{i=1}^{n}\left[\frac{x_i}{\mu(\mathbf{x})}\right]^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}}. \tag{24}$$

[7] Note that monotonocity is unneccesarily strong for this step: for example one could define ξ in cases where one required only that W is increasing if *all* incomes are increased by δ not just if *some* income is increased by δ. However, the assumption of monotonicity is useful for other results that follow.

[8] The limiting forms of (23) and (24) as $\varepsilon \to 1$ are, respectively, $\xi(\mathbf{x}) = \exp\left(\frac{1}{n}\sum_{i=1}^{n}\log(x_i)\right)$ and $I_{\mathrm{A}}^{1}(\mathbf{x}) = 1 - \exp\left(\frac{1}{n}\sum_{i=1}^{n}\log(x_i)\right)/\mu(\mathbf{x})$.

The number ε – the degree of (relative) inequality aversion – is a parameter that characterises individual members of the class of inequality measures. For any given unequal income distribution, the larger is ε the larger is the Atkinson inequality index – there is an example of this in Table 3 below. There is a close analogy with a class of risk indices in the case of constant relative risk aversion. This is unsurprising since this approach was explicitly founded on the formal similarity between distributional comparisons in terms of inequality and of risk (Atkinson 1970).

If, instead of scale-independence property, we required I to satisfy translation independence (14) then we would obtain a different class of indices

$$I_{\mathrm{K}}^{\beta}(\mathbf{x}) := \frac{1}{\beta} \log \left(\frac{1}{n} \sum_{i=1}^{n} e^{\beta[x_i - \mu(\mathbf{x})]} \right) \tag{25}$$

where $\beta > 0$ is a sensitivity parameter indexing members of the class (Kolm 1976). The connection of (25) with constant absolute risk aversion is evident.[9]

3.3 Ranking distributions

As noted earlier there are important results available about welfare and inequality comparisons that do not require the usage of specific indices. They follow from standard first- and second-order dominance results that are familiar from finance and other disciplines. Take the special class of *additive* welfare functions where W in (15) can be written in the form $\sum_{i=1}^{n} u(x_i)$ for some function u. If W is additive and satisfies the monotonicity axiom then u must be a strictly increasing function; if, furthermore, W satisfies the principle of transfers then u must be strictly concave. Then the following powerful results are available for any two distributions $\mathbf{x}'$ and $\mathbf{x}'' \in X$:

- *First-order.* $\mathbf{x}'$ strictly Parade-dominates $\mathbf{x}''$ if and only if $W(\mathbf{x}') > W(\mathbf{x}'')$ for any additive W that satisfies the principle of monotonicity.
- *Second-order.* $\mathbf{x}'$ strictly GLC-dominates $\mathbf{x}''$ if and only if $W(\mathbf{x}') > W(\mathbf{x}'')$ for any additive W that satisfies monotonicity and the principle of transfers (Shorrocks 1983).

A version of the second-order result applies to the conventional Lorenz curve and it accords with the intuitive argument presented in the introduction. Take the class of SWFs that satisfy the principle of transfers (they do not have to be additive). Then, for two distributions $\mathbf{x}'$ and $\mathbf{x}''$ that have the same mean, the statement "$W(\mathbf{x}') > W(\mathbf{x}'')$ for any W in this class" is true if and only if $\mathbf{x}'$ strictly Lorenz-dominates $\mathbf{x}''$. Furthermore, under these circumstances for any inequality index I that satisfies the principle of transfers it must be the case that $I(\mathbf{x}') < I(\mathbf{x}'')$. The implication of this is that all inequality measures

[9] One could also use "regularity" assumptions other than scale- or translation-independence – see Bossert and Pfingsten (1990).

that satisfy the principle of transfers "go the same way" if one distribution Lorenz-dominates the other. This is illustrated in Table 3 (again using the distribution of household income by households). Rows 1 to 4 give the results for the Atkinson indices: notice that in each case measured inequality is closer to 1 (the maximum) the higher is the degree of inequality aversion. The indices in the last two rows of Table 3 are discussed in the next section.

	1974	**2004**
$I_{\mathrm{A}}^{0.25}$	0.067	0.097
$I_{\mathrm{A}}^{0.5}$	0.134	0.190
$I_{\mathrm{A}}^{0.75}$	0.207	0.286
$I_{\mathrm{A}}^{1.0}$	0.297	0.418
I_{Gini}	0.395	0.466
I_{GE}^{0}	0.352	0.542
I_{GE}^{1}	0.267	0.406

Table 3: Inequality indices for the example in Table 1

4 Decomposition

The axioms discussed in section 2.2 induced some structure on inequality measures. By introducing the idea of *decomposing* inequality we can impose more structure and thereby obtain a useful class of indices. There are two principal types of decomposition: by subgroups of the population (regions, age groups,...) and by components of income (labour income, income from capital,...). Here we focus just on the population-subgroup issue.

Imagine that the population of n persons can be partitioned into a collection of m groups so that any individual falls into just one of these m groups. Each group j could be considered as a sub-population of size n_j in its own right (where $\sum_{j=1}^{m} n_j = n$) and one could compute inequality within this subpopulation as

$$\iota_j = I\left(\mathbf{x}_j\right) \tag{26}$$

where $\mathbf{x}_j$ is the income distribution consisting of just the members of subgroup j. The essence of the decomposition problem is to represent inequality overall as a function of inequality in each group $j = 1, ..., m$

$$I\left(\mathbf{x}\right) = F\left(\iota_1, \iota_2, ..., \iota_m; \pi_1, ..., \pi_m, s_1, ..., s_m\right) \tag{27}$$

where F is an aggregation function and the terms after the ";" show that aggregation may depend on the groups' shares of the population $\pi_j := n_j/n$ and the groups' shares of total income $s_j := n_j\mu\left(\mathbf{x}_j\right)/n\mu\left(\mathbf{x}\right)$. A consistency requirement on (27) is that if the income distribution within subgroup j changes so as

to increase ι_j in (26), all other things remaining the same, then inequality overall should increase. Insisting on this requirement on F for all logically possible partitions induces a type of separability on the function $I(\cdot)$ so that the index must be of the general form mentioned in footnote 6 above. If we also require that scale-independence hold then the inequality index must take the specific form

$$I^{\alpha}_{\mathrm{GE}}(\mathbf{x}) = \frac{1}{\alpha^2 - \alpha}\left[\frac{1}{n}\sum_{i=1}^{n}\left[\frac{x_i}{\mu(\mathbf{x})}\right]^{\alpha} - 1\right] \tag{28}$$

or some monotonic transform of it, where α is a real number. The "GE" used in the labelling of (28) stands for the *generalised entropy* class, which is a generalisation of the two indices introduced by Theil (1967).[10] The α in (28) is a parameter that characterises different members of the GE class: a high positive value of α yields an index that is very sensitive to income transfers at the top of the distribution; specifying a negative value will produce an index that is sensitive to income transfers among the poor.[11]

5 Income differences

The third way forward from the basic argument outlined in section 2.2 focuses on fundamental income differences. This is one of the key ways in which one can motivate usage of the very well-known inequality indices mentioned in section 2.1.2. The variance and the coefficient of variation (4) can be thought of as a representation of the averaged squared difference between each income x_i and the mean. A compelling argument for the Gini coefficient is that it is the (normalised) expected value of the absolute difference between any two randomly selected incomes in the population.

However, there are other types of income difference that are of special relevance to inequality measurement. Just as some poverty indices can be characterised as a kind of average distance of individual incomes from a reference income level, the poverty line,[12] so also some inequality measures use the idea of a reference level income. In the case of inequality the reference level income level has been suggested as either that of the best-off person in society, or the average income of all those who are better off than any given person i (Temkin 1993). In each of these cases application of standard axioms about the structure of inequality orderings leads to a class of inequality indices that bears a functional similarity to poverty indices and to indices of relative deprivation (Cowell and Ebert 2004).

[10] Theil's two indices are those corresponding to the special forms in the cases $\alpha = 0, 1$: $I^{0}_{\mathrm{GE}}(\mathbf{x}) := -\frac{1}{n}\sum_{i=1}^{n}\log\left(x_i/\mu(\mathbf{x})\right)$ and $I^{1}_{\mathrm{GE}}(\mathbf{x}) = \frac{1}{n}\sum_{i=1}^{n}\left[x_i/\mu(\mathbf{x})\right]\log\left(x_i/\mu(\mathbf{x})\right)$. The values of these indices for the US example are given in the last two rows of Table 3.

[11] There is a functional relationship between the the class (24) and the class (28). For any $\alpha < 1$ we have $I^{\varepsilon}_{\mathrm{A}}(\mathbf{x}) = 1 - \left[1 + \alpha\left[\alpha - 1\right] I^{\alpha}_{\mathrm{GE}}(\mathbf{x})\right]^{1/\alpha}$ where $\varepsilon = 1 - \alpha$.

[12] Many poverty indices can be written in the form $\frac{1}{n}\sum_{i=1}^{n} p\left(z - x_i\right)$ where z is the poverty line and $p(\cdot)$ is a nondecreasing function that is zero for all $x_i \geq z$.

6 Implementation

The practical issues associated with the exposition of the example in Tables 1 and 2 highlight some of the problems in implementing inequality measures and associated tools – the definition of income, income receiver and so on. Given the way in which income data are usually obtained, issues of sampling and measurement error usually need to be treated carefully. Furthermore, the special nature of income and wealth distributions and the sensitivity of inequality indices to very high or very low incomes usually requires that particular attention be given to the problem of outliers. Finally it should be noted that it is still sometimes the case that the data required for estimating inequality indices are only made available in grouped form rather than as microdata so that special techniques may be required for interpolation within income intervals and for modelling the tails of the distribution.

7 Further reading

For the welfare-economic issues see Atkinson (1983) and Sen and Foster (1997). For literature surveys see Cowell (2000, 2007) and Lambert (2001).

References

Atkinson, A. B. (1970). On the measurement of inequality. *Journal of Economic Theory 2*, 244–263.

Atkinson, A. B. (1983). *The Economics of Inequality* (Second ed.). Oxford: Clarendon Press.

Bossert, W. and A. Pfingsten (1990). Intermediate inequality: concepts, indices and welfare implications. *Mathematical Social Science 19*, 117–134.

Cowell, F. A. (2000). Measurement of inequality. In A. B. Atkinson and F. Bourguignon (Eds.), *Handbook of Income Distribution*, Chapter 2, pp. 87–166. Amsterdam: North Holland.

Cowell, F. A. (2007). *Measuring Inequality* (Third ed.). Hemel Hempstead: Oxford University Press.

Cowell, F. A. and U. Ebert (2004). Complaints and inequality. *Social Choice and Welfare 23*, 71–89.

Dalton, H. (1920). Measurement of the inequality of incomes. *The Economic Journal 30*, 348–361.

DeNavas-Walt, C., B. D. Proctor, and C. H. Lee (2005). Income, poverty, and health insurance coverage in the United States: 2004. Current Population Reports P60-229, U.S. Census Bureau, U.S. Government Printing Office, Washington, DC.

Kolm, S.-C. (1976). Unequal inequalities I. *Journal of Economic Theory 12*, 416–442.

Lambert, P. J. (2001). *The Distribution and Redistribution of Income* (Third ed.). Manchester: Manchester University Press.

Pen, J. (1974). *Income Distribution* (second ed.)., Chapter 3, pp. 48–59. London: Allen Lane, The Penguin Press.

Sen, A. K. and J. E. Foster (1997). *On Economic Inequality* (Second ed.). Oxford: Clarendon Press.

Shorrocks, A. F. (1983). Ranking income distributions. *Economica 50*, 3–17.

Temkin, L. S. (1993). *Inequality*. Oxford: Oxford University Press.

Theil, H. (1967). *Economics and Information Theory*. Amsterdam: North Holland.