# Income Inequality, Population Heterogeneity, and Mobility: Some Basics

John Creedy

WORKING PAPER 02/2023
April 2023

Working Papers in Public Finance

VICTORIA UNIVERSITY OF WELLINGTON
TE HERENGA WAKA

Chair in
Public Finance

**The Working Papers in Public Finance** series is published by the Victoria Business School to disseminate initial research on public finance topics, from economists, accountants, finance, law and tax specialists, to a wider audience. Any opinions and views expressed in these papers are those of the author(s). They should not be attributed to Victoria University of Wellington or the sponsors of the Chair in Public Finance.

Further enquiries to:
The Administrator
Chair in Public Finance
Victoria University of Wellington
PO Box 600
Wellington 6041
New Zealand

Phone: +64-4-463-9656
Email: cpf-info@vuw.ac.nz

Papers in the series can be downloaded from the following website:
https://www.wgtn.ac.nz/sacl/centres-and-chairs/cpf/publications/working-papers

# Income Inequality, Population Heterogeneity, and Mobility: Some Basics*

John Creedy
Victoria University of Wellington

**Abstract**

This paper explores some basics of income inequality measurement, and challenges arising when extending the accounting period. When discussing inequality measurement, the emphasis is on basic conceptual aspects, particularly the role played by explicit value judgements. In considering incomes over several periods, a greater emphasis is given to more technical descriptive methods of analysis. In assessing distributions, a value judgement that often plays a central role is that of anonymity. However, this value judgement must be rejected in (at least) two important cases. First, many individuals belong to families or households, within which some form of income sharing takes place. Second, when incomes fluctuate over time, it is necessary to be able to trace the precise income paths followed by different individuals, who may of course also move between family or household types. The paper emphasises the importance of these two aspects of heterogeneity.

*This paper is part of a project on 'Measuring Income Inequality, Poverty, and Mobility in New Zealand', funded by an Endeavour Research Grant from the Ministry of Business, Innovation and Employment (MBIE)\ and awarded to the Chair in Public Finance at Victoria University of Wellington. In writing this paper, I have benefited from discussions with Subbu Subramanian. The sections on mobility diagrams draw on joint work with Norman Gemmell.

# 1 Introduction

The aim of this paper is to explore some basic aspects of income inequality measurement, and to examine some of the challenges involved when extending the accounting period to more than, say, a single year. The term 'income' is used throughout for convenience, but of course the analyses apply to any chosen 'welfare metric'. The paper is written at several 'levels'. When discussing inequality measurement, the emphasis is on basic conceptual aspects which have been extensively discussed in a huge literature over the last fifty years. Of importance is the way in which the implications of adopting alternative explicit value judgements can be examined. But in considering incomes over several periods, a greater emphasis is given to more technical descriptive methods of analysis. In some cases, numerical examples are provided to illustrate various properties, based on simulated data.

In assessing distributions, a basic value judgement that often plays a central role is that of anonymity (also referred to as 'symmetry'), such that the non-income characteristics and identities of individual income earners are not considered to be relevant. However, this value judgement must be rejected in (at least) two important cases. First, despite the fact that the 'individual' is in some sense a fundamental unit of analysis, it is recognised that many individuals belong to families or households, within which some form of income sharing takes place. In comparing the incomes of two individuals, it therefore becomes important to know the family or household to which those individuals belong: they cannot be treated as anonymous. Second, when it is desired to consider inequality among individuals whose incomes fluctuate over time, it is necessary to be able to trace the precise income paths followed by different individuals, who may of course also move between family or household types. Again, individual income profiles rely on being able to identify those individuals: they cannot be anonymous. This does not mean that their names and addresses must be known: some form of identifier is required so that family members can be linked, and individual incomes can be linked over time.

This paper arose partly from frustration with popular debates, where the level of discussion rarely rises above rhetoric. References are made to 'increasing inequality' but there is often no clarity about what is considered to be unequally distributed, or even what the population group is taken to be. Empirical summary measures are seldom reported, and when some quantitative indices are produced (nearly always Gini measures), they are used merely as passing illustrations, without any clarifying discussion. Furthermore, attention is typically restricted to the inequality of some flow measure, such as income, measured over a single year. The fact that individuals' incomes vary substantially from year to year is seldom given the recognition it deserves.

While aiming to raise an awareness of some of the basic analytics involved, no attempt is made to provide a comprehensive, or even representative, treatment. That is far too great a challenge. It merely attempts to explain some of the difficulties faced, and approaches that can be taken when summarising inequality and mobility. It is a highly personal (or indeed even parochial and idiosyncratic) selection, with no pretence to balance in selecting topics. No attempt is made to provide an extensive or balanced list of references to the enormous literature on the topics covered. However, references can be found in the papers cited.[1]

Section 2 discusses some very simple descriptions of income distributions, including the famous Lorenz curve and the inequality measure most closely associated with it, the Gini measure (although it was not initially introduced with the Lorenz curve in mind). Using the Lorenz curve, unambiguous judgements about the extent of inequality, and changes in it over time or between samples, can sometimes be made. The Gini measure enables all the many individual values to be compressed into a meaningful single summary measure. Section 3 turns to the way in which value judgements can be introduced in deriving inequality measures and making distributional comparisons. Emphasis is given to the Gini and Atkinson inequality measures, which are both seen to belong to a special class of indices. The approach is one of investigating the implications of adopting particular explicit value judgements, rather than investigators imposing their own values. It is seen that any attempt to provide a quantitative indication of inequality and mobility involves very strong assumptions. But the need to make strong assumptions need not lead to a nihilistic attitude. There is certainly a need for evidence, and an understanding of its limitations, to replace rhetoric.

As mentioned above, the basic analyses of inequality measures assume a population of individuals who are essentially homogeneous. That is, the value judgement is made that no non-income characteristics are relevant: all individuals are anonymous. Section 4 explores problems that arise where there are non-income differences among individuals arising from the fact that typically many individuals are part of a family or household, within which some income sharing is assumed to take place.

Section 5 then explores another important context in which anonymity is not appropriate, which arises when looking at income distributions in different time periods: this concerns the relationship between income mobility and inequality. For some purposes the mobility process itself may be of special concern, while in other contexts

[1] Bibiographies and more extensive analyses can be found in the following papers by the author. On Atkinson and Gini measures in Sections 2 and 3, see Creedy (2016, 2023a, b). For Section 4, see Creedy and Sleeman (2005) and Creedy (2017). On income mobility in Section 6, see Creedy (1975, 1977, 1985, 1997a, b). Sections 7 to 11 draw on Creedy and Gemmell (2018, 2019, 2022, 2023).

attention may be focused on an income measure obtained using a longer accounting period (such as the present value of an income stream). Section 6 uses simple dynamic models to explore inequality changes over time, and the use of different accounting periods. It is seen that, even in such simple models, the relationships among different distributions are far from clear. The relationship between longer-period measures and cross-sectional distributions, when there are overlapping cohorts, is also explored.

Section 7 discusses some simple ways that have been used to illustrate the nature of differential income changes. Section 8 turns to a more recent diagrammatic approach that simultaneously reflects three elements of income mobility: these are incidence, intensity and inequality. For this reason the diagrams are referred to as 'Three "I"s of Mobility, or TIM, curves. Section 9 presents a summary measure of equalising mobility, based on the TIM curve. Section 10 turns to the case where concern is with a particular group of low-income earners, and is this concerned with transitions across a threshold low-income, or 'poverty', line. Section 11 turns from mobility seen in terms of relative income changes to mobility viewed in terms of changes in the rank-order of people in the distribution. Brief conclusions are in Section 12

# 2 Simple Descriptions of Distributions

This introductory section discusses a diagrammatic approach to comparing distributions that plays a central role. First, Subsection 2.1 shows that standard representations of distributions are inadequate. The most popular form of diagram, the Lorenz Curve, is described in Subsection 2.2, and the associated most popular summary measure of inequality, the Gini Measure, is described in Subsection 2.3.

## 2.1 Basic Diagrams

Consider the question of how to describe the form of an income distribution, containing a large number of observations from a specified population group at a given date, such that comparisons can be made with other distributions. For illustrative purposes, a simulated sample of 500 incomes was generated by taking random draws from a particular functional form of the income distribution, the Lognormal. This form, which is known to provide an approximation to many positively skewed distributions of essentially positive variables, assumes that the logarithms of income follow the famous Normal distribution. The arithmetic mean and variance of log-income were set arbitrarily at 3 and 0.5. A simple plot of the income values, after arranging incomes in ascending order, gives very little indication of the nature of inequality, as seen from Figure 1.

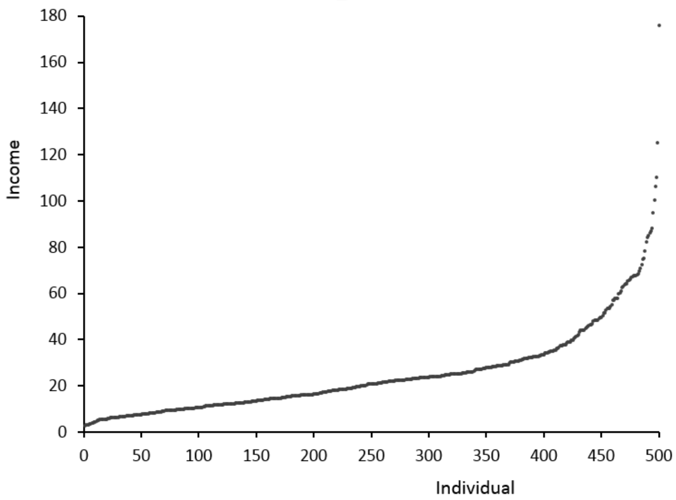


Figure 1: Simple Plot of Incomes in Ascending Order

A much better indication is given by plotting a histogram, showing the frequency (or density) function. This provides a valuable picture of the whole range of incomes. The histogram can reveal the nature of any skewness and peakedness (or kurtosis), and can highlight the possible existence of multimodality. In particular, any bunching of individuals at certain income levels or ranges is revealed, perhaps associated with income tax thresholds, or social benefits, depending on the income measure used. However, the shape can be distorted to some extent by the choice of income classes (which can influence the smoothness of the histogram). Examples are shown in Figure 2, for the simulated data described above. Even though the random draws are from a smooth, regular, distribution, the histograms show a jagged profile. The fact that the view is influenced by the class widths chosen to group individual observations is shown by comparing the two histograms.

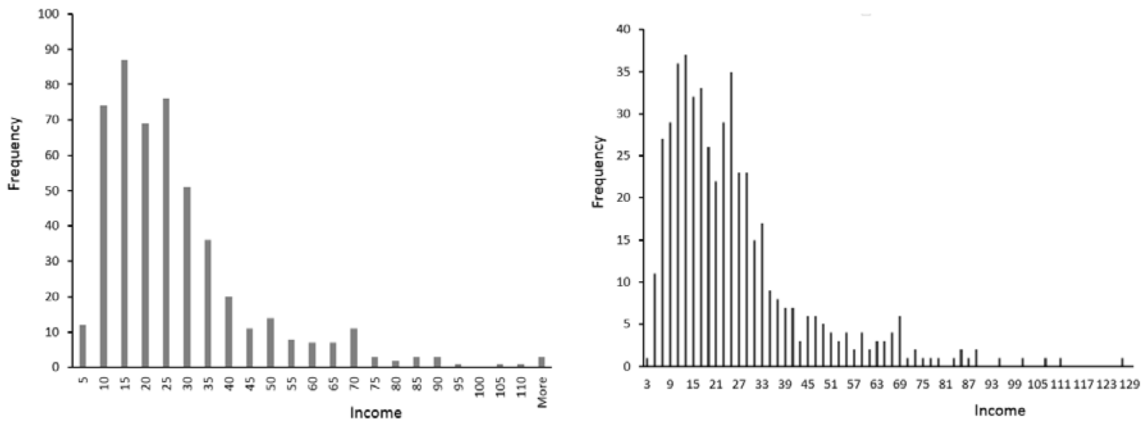


Figure 2: Two Histograms Using The Same Data

In order to make further comparisons, summary measures of such a large complex picture are required. Where income levels are concerned, standard measures of location, such as the arithmetic or geometric mean, or median, are used. But care must be taken, because two distributions may have quite different mean values, and yet there can be a considerable overlap: the distribution with the lower mean may nevertheless contain many individuals with incomes higher than a large number of people belonging to the distribution with the higher mean. It is even possible for the arithmetic means and medians to give different rankings of two distributions, depending on their kurtosis and skewness.

## 2.2 The Lorenz Curve

Regarding visualisations of distributions, where the dispersion of incomes is of interest, the most famous diagram is without question the Lorenz Curve. This is defined for non-negative values, and named after the American economist Max Lorenz, who first devised the diagram in an undergraduate essay, and published it while he was a graduate student. This shows, after placing the incomes in ascending order, the relationship between the cumulative proportion of total income and the corresponding proportion of individuals. Since incomes are in ascending order, the curve must be below the diagonal. Equal incomes produce a upward sloping diagonal line running from the bottom left to the top right corner of a box with unit sides. For a large sample size, the extreme inequality case is represented by the bottom and right hand sides of the box. The Lorenz curve for the data used in the above histograms is shown in Figure 3. Given the use of population proportions and income shares, the curve is independent of the units in which income is measured. Also, adding an identical population to an existing one has no effect on the Lorenz curve.

The diagram also shows that, being based on cumulative values, the Lorenz curve (just like the cumulative distribution, or ogive, or distribution function) is 'smoothed out'. The jagged histogram is converted to a relatively smooth curve. The symmetry of the curve in this diagram arises because the income values were generated as random observations from a Lognormal distribution, for which Lorenz curves are necessarily symmetric. But of course it is not a general property.

Two additional features are worth stressing. First, it is clear that a transfer of income from a richer to a poorer person, with no change in their rank order, must move the Lorenz curve closer to the diagonal line of equality, given the cumulation process involved. Second, if the Lorenz curve for one distribution, say A, lies everywhere closer to the line of equality compared with another distribution, B, it is said that A 'Lorenz dominates' B. If Lorenz curves intersect, no such unambiguous comparison can be

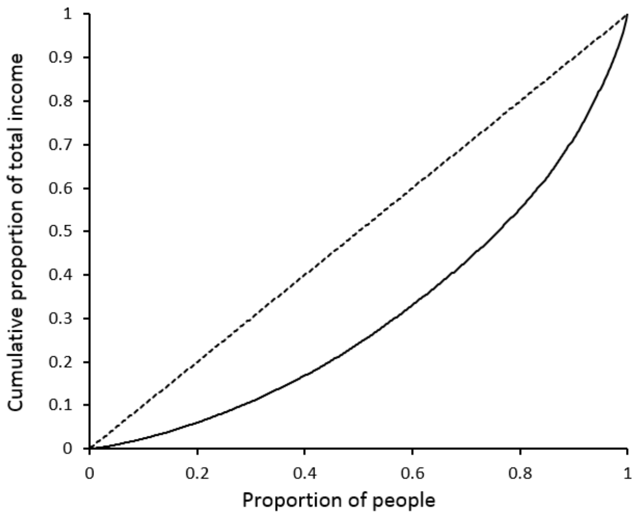


Figure 3: Lorenz Curve for Simulated Data

made, without a summary measure and associated assumptions.

## 2.3 The Gini Measure and the Lorenz Curve

The Lorenz curve suggests that an inequality measure can be based on the area between the Lorenz curve and the diagonal line corresponding to equal incomes, as a ratio of the maximum inequality case (for which the area is one half). It is a normalised 'area measure' of the 'distance' of the Lorenz curve from the line of equality. The Gini coefficient thus lies between 0 and 1. It turns out that this area is the same as the relative mean-difference between all pairs of incomes, which was first proposed, in the early 1900s, as an inequality measure by the famous Italian statistician, Gini. A number of different expressions can be found for the Gini measure, $G$, but one of the most frequently used, for ungrouped data, is as follows. First, the incomes are arranged in non-descending order such that $x_1 \leqslant x_2 \leqslant .... \leqslant x_n$. Then:

$$G = \frac{n+1}{n} - \frac{2}{n^2 \bar{x}} \sum_{i=1}^{n} (n+1-i)\, x_i \tag{1}$$

Using (1), the Gini inequality measure for the simulated data is found to be 0.381.

Other statistical indices of inequality include simple comparisons of measures of location, such as the inter-quartile range, or the ratio of the 90th to the 10th percentiles, although these necessarily ignore many actual values. Other basic statistical measures involve the use of 'moments', in particular the variance (the second moment about the arithmetic mean), or its square root, the standard deviation, or the latter's normalised value, the coefficient of variation. The statistical properties of these measures, particularly the form of their sampling distributions, are well established. One alternative route is to fit a functional form to the data, from which a wide range of measures can be obtained, given the moment generating function of the specified distribution and the estimated parameter values. However, the imposition of a smooth functional form is in many cases not appropriate, unless it is needed in other economic models or where data are very limited (for example with grouped distributions and wide income classes).

It is worth recognising a property of any summary measure: a wide range of distributions may be consistent with the same numerical value of the measure. In the case of the Gini measure this is easily illustrated by the fact that it is possible to draw a number of different Lorenz curves, for which the area between each curve and the line of equality is in each case the same.

# 3 Introducing Value Judgements

A feature of statistical measures is that they make no reference to value judgements. Two basic approaches have been used. One is to draw up a list of properties, or statements of value judgements, that the investigator wants an inequality measure to have: these properties are referred to as axioms. Any statistical measure can then be investigated to see if it conforms to those axioms. Of course, it is an open question as to how widely the various axioms typically used are shared more widely, and this has given rise to a literature of its own, based on various types of survey. One of the most important axioms is the 'principle of transfers', according to which an income transfer from a richer to a poorer individual, which preserves their ranks, is judged to represent an improvement. For example, from the previous section, it is immediately clear that the Gini measure satisfies this principle.

Other axioms widely used include the following (loosely stated). 'Scale invariance' states that the measure is invariant with respect to equal proportional changes in all incomes, while 'translation invariance' relates to equal absolute incomes changes. 'Population invariance' states that the measure is invariant with respect to replications of the population. 'Anonymity' (or 'Symmetry') requires that the measure does not identify individuals: for example, a three-person population with incomes 4, 2, and 5 for persons 1, 2 and 3 respectively has the same measure as the distribution with the same individuals and incomes of 2, 5 and 4. The Gini therefore satisfies the principle of transfers, scale invariance, population invariance and anonymity. Another axiom sometimes applied requires the measure to be 'subgroup consistent', whereby if the population is divided into groups, an overall measure can be decomposed such that an increase in inequality for any group produces an increase for the population as a whole.

The other approach to introducing value judgements is to derive a measure directly from an explicit form of evaluation function. This evaluation function is referred to as a Social Welfare Function, or SWF, which, despite its name, does not refer to 'society' but is taken to represent the value judgements of a hypothetical independent disinterested judge. This is the approach explored further in this section. Subsection 3.1 begins by introducing a class of measures, starting from a given SWF. Subsection 3.2 describes an eponymous measure belonging to this class, proposed by Atkinson. Subsection 3.3 shows that the Gini measure can also be derived as a member of the same broad class of inequality measures.

## 3.1 A Class of Measures

A widely-adopted approach is to use a member of a broad class of measures defined in terms of the proportional difference between the arithmetic mean income and the equally distributed equivalent income. The latter is defined as that income which, if obtained by everyone, gives the same value of the social welfare function as the actual distribution. The measure therefore involves a comparison of two measures of location. An implication of the approach is that a judge who is concerned only with aggregate income, and has no aversion to inequality, arrives at a measure of 'inequality' equal to zero, even though the statistical dispersion may be large. In that case the equally distributed equivalent income is equal to the arithmetic mean.

For incomes of $x_i$, for $i = 1, ..., n$, denote the arithmetic mean and equally distributed equivalent as $\bar{x}$ and $x_{EDE}$ respectively. The social welfare function can be written in general as $W = W(x_1, ..., x_n)$. Hence, an expression for $x_{EDE}$ can be obtained by solving:

$$W(x_{EDE}, ..., x_{EDE}) = W(x_1, ..., x_n) \tag{2}$$

Inequality, $I$, is:

$$I = 1 - \frac{x_{EDE}}{\bar{x}} \tag{3}$$

The welfare function is regarded as encapsulating the distributional value judgements – the nature of aversion or otherwise to inequality – of an independent disinterested judge. The choice of measure thus depends on the choice of welfare function, defined in terms of incomes, and this is often a type of weighted sum. Importantly, the 'principle of transfers' discussed above simply requires the $W$ function to be concave. This means that the reduction in $W$, arising from a given reduction in a higher income, is less than the increase in $W$ that arises from the same absolute increase in a lower income. These judgements may of course reflect, or subsume, a number of the axioms mentioned above. It is worth stressing again that the social welfare function is purely an evaluation function of a fictitious judge. Individual incomes in some way contribute to total 'social welfare' and, despite the language used, a single judge is involved, rather than a society or other aggregate.

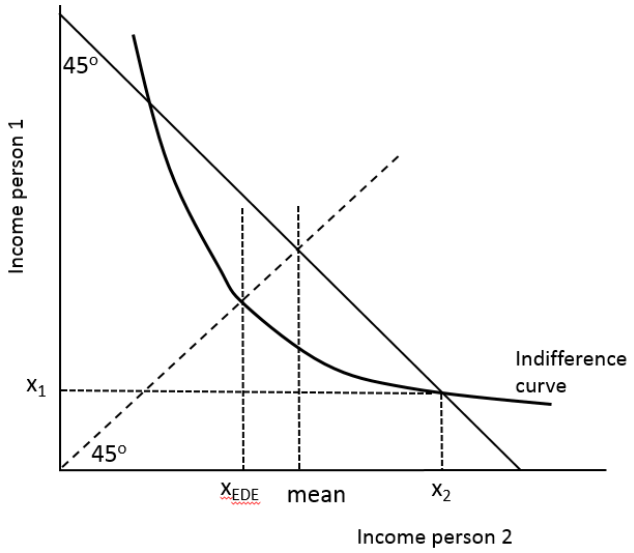


Figure 4: A Social Indifference Curve and Equally Distributed Equivalent Income

The approach can be illustrated using a convenient diagram, shown in Figure 4. The two axes measure the incomes of two individuals, 1 and 2. The location of the two individuals is given by the coordinates, incomes $x_1$ and $x_2$. Given $W$, it is possible to draw a social indifference curve, showing the locus of allocations that produce the same value of $W$ as the actual distribution. The convexity of this curve reflects the principle of transfers. The symmetry of the indifference curve around the upward-sloping 45 degree line from the origin (along which incomes are equal) reflects the anonymity property. The intersection of this 45 degree line with the backward-sloping 45 degree line through the actual allocation gives the arithmetic mean, $\bar{x}$. This backward-sloping line corresponds to the indifference curve that would result from the complete absence of any aversion to inequality. The intersection of the indifference curve with the line of equality gives the equally distributed equivalent, $x_{EDE}$.

This view of inequality is thus one in which the actual allocative mechanism is seen as generating some 'inefficiency'. That is, a higher value of social welfare, $W$, can be achieved with a perfectly equal distribution, of the same total income, compared with the actual distribution. Perfect equality would produce $W(\bar{x}, ..., \bar{x})$ whereas the observed distribution produces $W(x_1, ..., x_n)$, which by definition is the same as $W(x_{EDE}, ..., x_{EDE})$. Hence the proportional difference between $\bar{x}$ and $x_{EDE}$ can be said to reflect the allocative inefficiency – here termed inequality. This perspective helps to clarify the apparent paradox that, if the judge has no aversion to inequality, a distribution that has some statistical dispersion can be referred to as having zero inequality.

### 3.2 The Atkinson Inequality Measure

A widely used welfare function, following Atkinson, is to write $W = \sum_{i=1}^{n} \Phi(x_i)$, with $\Phi(x_i)$ taking the form, for $\varepsilon \neq 1$:

$$\Phi(x_i) = \frac{x_i^{1-\varepsilon}}{1-\varepsilon} \tag{4}$$

and for $\varepsilon = 1$, $\Phi(x_i) = \log(x_i)$. This reflects a constant relative aversion to inequality, of $\varepsilon = x_i \Phi''(x_i) / \Phi'(x_i)$. Then $x_{EDE}$ can be solved as:

$$x_{EDE} = \left[\frac{1}{n}\sum_{i=1}^{n} x_i^{1-\varepsilon}\right]^{1/(1-\varepsilon)} \tag{5}$$

from which Atkinson's measure, $I_A$, is directly obtained using (3). In terms of Figure 4, it can be shown that the social indifference curves arising from this welfare function are homothetic: the slopes are constant along a ray from the origin.

Unlike the Gini measure, Atkinson's index does not depend on rank-order positions of individuals. It is a simple matter to examine the implications of adopting any specified degree of relative inequality aversion, and to show that the rankings of a range of distributions by inequality can change substantially as $\varepsilon$ is varied. Using the simulated data described earlier, $I_A$ takes values of 0.048, 0.116 and 0.390 respectively for inequality aversion parameters of 0.2, 0.5, and 2.0. Measured inequality obviously increases as $\varepsilon$ increases, although the precise way in which it increases is data-specific.

An infinitely large $\varepsilon$ produces an equally distributed equivalent income equal to the minimum income: if this is very small in relation to the arithmetic mean, the inequality measure is close to 1. The judge with such an extreme aversion discounts (sets to zero) all incomes in excess of the minimum income: it is thus referred to as a maxi-min rule. In trying to get some idea of what different values of $\varepsilon$ mean for inequality aversion, Atkinson proposed adopting the 'thought experiment' of the 'leaky bucket'. That is, if a given amount is taken from a rich person, what is the minimum amount that must be given to a specified poorer person for $W$ to remain unchanged? The difference between the two amounts represents the maximum leak that the judge would tolerate in making a transfer. From 4, the slope of a social indifference curve through the point, $[x_j, x_k]$ is:

$$\left.\frac{dx_j}{dx_k}\right|_W = -\left(\frac{x_k}{x_j}\right)^{-\varepsilon} \tag{6}$$

Hence, if $x_k > x_j$, and converting to small discrete changes, suppose \$1 is taken from person $k$, so that $\Delta x_k = -1$. The amount that must be given to person $j$ for $W$ to remain unchanged is thus:

$$\Delta x_k = \left(\frac{x_k}{x_j}\right)^{-\varepsilon} \tag{7}$$

Suppose $k$ has twice as much as $j$, and $\varepsilon = 1$. The amount given to $j$ is therefore 50 cents. Equation (7) can be used to form an appreciation of what different values of $\varepsilon$ mean in terms of inequality aversion.

An implication of the this type of concave function is that the inequality measure is less sensitive to inequality-reducing income transfers at the top of the distribution compared with transfers in the lower-income ranges. The welfare function is said to satisfy a further condition regarding transfers: the 'principle of diminishing transfers' implies that a 'progressive' transfer at the lower end of the distribution should have a larger effect on measured inequality than a transfer in the higher-income ranges.

In the simulated sample suppose the income of the richest person is increased by a multiple of four. This places the richest person some considerable distance from other members of the sample. The Gini measure increases to 0.404, while the Atkinson measures (for the three values of $\varepsilon$ used above) increase to 0.063, 0.142 and 0.415. The increase is 31 per cent for the case of $\varepsilon = 0.2$, but is only 6 per cent for the higher aversion coefficient of 2.0. This means that any special focus on top incomes, in debates about income inequality, implies distributional value judgements that do not agree with the form used in the Atkinson measure.

However, it is possible to specify an alternative welfare function which reflect both the principle of transfers and the value judgement that an inequality-reducing transfer should have a larger effect at the higher end of the distribution, compared with similar transfers at the lower end. An extreme form of this kind of aversion produces the prescription of minimising the maximum income (remembering that the rank order must remain unchanged): this is a mini-max rule, contrasting with the maxi-min rule discussed above.

### 3.3 A Welfare Function and the Gini Measure

It can be shown that the class of inequality measures discussed in Subsection 3.1 also includes the Gini measure. The associated welfare function attaches weights to each income that depend inversely on their rank-order positions, with incomes arranged in ascending order. Hence, with $x_1 \leq x_2 \leq ... \leq x_n$, $W$ takes the form:

$$W_B = \sum_{i=1}^{n} (n + 1 - i)\, x_i \tag{8}$$

The resulting equally distributed equivalent income, $x_{E,G}$ is given by:

$$x_{E,G} = \sum_{i=1}^{n} \left\{ \frac{n + 1 - i}{\sum_{j=1}^{n} (n + 1 - j)} \right\} x_i \tag{9}$$

This gives the corresponding Gini-type inequality measure as:

$$I_G = 1 - \frac{x_{E,G}}{\bar{x}} \tag{10}$$

Close comparison of (10) with (1) reveals that they differ slightly for small $n$. However, for all practical cases, with large $n$, they give identical values. Exploration of (8) quickly shows that the nature of the willingness to tolerate a leaky bucket in making transfers between two individuals (with their ranks unchanged) differs substantially from the Atkinson measure. Various extensions to this basic Gini measure are also available, using modifications of the form in (8).

# 4 Non-Homogeneous Individuals

The basic approach described above necessarily involves some strong assumptions. In particular, it assumes that the population consists of homogeneous individuals: essentially there are assumed to be no relevant non-income differences between the individuals. This assumption lies behind the anonymity property of the social welfare functions commonly used (and the associated symmetry of social indifference curves around the line of equality in Figure 4).

However, in practice there are many characteristics that become evident once people are asked to judge income differences, even among very few people. Asked to give a view about hypothetical income differences, people quickly raise questions concerning any differences in health, disability, labour market attachment, hours worked and the intensity of work. Other characteristics may include the precise source of income (employment, rental, interest, capital gains), occupation, differences in innate ability, previous human capital and other investments, differences in previous saving choices, and so on. As soon as the analysis departs from a context in which all individuals are homogeneous, it becomes clear that anonymity is no longer a desirable property of the welfare function.

In practice the major, indeed usually the only, non-income difference accommodated by the approach relates to the fact that individuals live in family or household units, which vary in terms of their size and composition. Each unit, consisting of more than one person, is reduced to a scalar size, the number of adult equivalent people in the unit, which is obtained by the application of a set of adult equivalent scales. The strong assumption is then made that within the relevant unit there is equal sharing among individuals (or that non-equal sharing is not considered relevant to the comparisons). The choice of unit raises numerous complications, but for convenience in what follows it is assumed that the 'family' (consisting of people who are related/partners and dependents living together) is the chosen unit. Suppose there are $N$ families, each of which has $n_i$ individuals and a total family income of $y_i$. Using a set of adult equivalent scales (to be considered further below), the adult equivalent size of family $i$ is denoted, $s_i$. The welfare metric is thus, $x_i = y_i/s_i$.

Consider the approach proposed by Atkinson, discussed above. The appropriate welfare function is thus a weighted sum of $N$ terms, $\Phi(x_i)$, with $x_i = y_i/s_i$. The important next step involves making a decision about the weights to be used. A number of empirical studies take the family itself as the basic unit, implying a welfare function of the form, $W_H = (1/N)\sum_{i=1}^{N}\Phi(x_i)$. However, it is hard to find any welfare rationale for such a choice. An alternative, widely used choice (too often left implicit), is to treat the individual as the basic unit of analysis. This involves using a welfare function as the weighted sum:

$$W_I = \sum_{i=1}^{N}\left(\frac{n_i}{\sum_{i=1}^{N} n_i}\right)\Phi(x_i) \tag{11}$$

The same weights are used in calculating $\bar{x}$ and $x_{EDE}$. Each individual is assigned the welfare metric of the relevant family, and effectively 'counts for one' irrespective both of the family to which the individual belongs and the person's adult equivalent size. This implies that inequality remains unchanged when one person (of whatever type) in the population is replaced by another person having the same welfare metric, $x$, but belonging to any other type of family. Importantly, this approach does not in general satisfy the principle of transfers. An income transfer from a poor to a richer (and larger) family can reduce inequality and raise social welfare. The reason for this paradox is that, with economies of scale in the choice of adult equivalent scales, large families are regarded as being 'more efficient' at generating social welfare. Hence, Lorenz dominance in this context does not necessarily imply welfare dominance: the equivalence of the two is a fundamental component of welfare analysis for homogeneous populations.

An alternative is to define the basic unit of analysis as the 'adult equivalent person'. This means that the associated welfare function is:

$$W_E = \sum_{i=1}^{N} \left( \frac{s_i}{\sum_{i=1}^{N} s_i} \right) \Phi(x_i) \tag{12}$$

This approach means that an individual's contribution to inequality and social welfare depends on the composition of the family to which that person belongs. For example, an adult in a one-person family 'counts for one', but the same person in a family containing another adult and and several children counts for 'less than one'. A feature of this approach is that it satisfies the principle of transfers: a transfer from a richer to a poorer 'equivalent person' results in a reduction in inequality (and an increase in social welfare). The choice between individuals and adult equivalents as the basic unit of analysis in inequality and social welfare calculations therefore involves a choice between two incompatible value judgements. They can in principle lead to opposite conclusions about the effects of a tax policy change on inequality.

Unfortunately, this complication is not as widely recognised as it should be, and empirical studies seldom make the choice of unit explicit. Furthermore, many empirical studies take a rather cavalier attitude to the choice of adult equivalent scales, often referring only to OECD scales, without defining them. However, judgements regarding comparisons over time, or among different tax and transfer systems, can depend on the choice of scales, as well as the choice of the unit of analysis. Sensitivity analyses can easily be carried out using parametric scales. Thus, define $n_{a,i}$ and $n_{c,i}$ respectively as the number of adults and children in family, $i$. Consider the scales:

$$s_i = (n_{a,i} + \theta n_{c,i})^{\alpha} \tag{13}$$

where $\theta$ is the weight attached to children and $\alpha$ is a term reflecting the extent of economies of scale within the family. This simple formulation actually provides, by suitable choice of parameters, a very close approximation to many scales used in practice.

# 5 Mobility and Inequality

Previous sections have concentrated on inequality in a single-period context. The recognition that individual incomes (along with family and household units) do not remain constant over time provides the motivation for this and the following sections. To set the scene, consider a population of just three individuals with incomes in the first period of $(10, 20, 30)$ and no relevant non-income characteristics. Suppose, in moving to the second period, the incomes become $(15, 20, 30)$: this represents a Pareto improvement, with only an addition to the first person's income of 5 units, and the change is clearly equalising. However, compare this with an alternative transition producing the vector, $(20, 15, 30)$. This has the same total income in the second period as with the first transition, although when individuals are arranged in ascending order, persons 1 and 2 change places. Based only on cross-sectional comparisons, and invoking anonymity, the two alternative distributions, $(15, 20, 30)$ and $(20, 15, 30)$, would be regarded as identical. However, there is clearly a case for abandoning anonymity: the population heterogeneity here concerns the nature of the income time streams, and the re-ranking that occurs over time is an important consideration. In this example, the two alternative transitions are different, but there is no reason to regard one as being more equalising than the other. Furthermore, if comparisons are made on the basis of alternative distributions of two-period incomes, it is easy to see that the two relevant Lorenz curves intersect.

Comparisons of different time streams, or patterns of mobility, are highly complex. Compared with the simple case of constant income streams, it would doubtless be considered desirable for individuals to have an opportunity to improve their situation by hard work, or investment in education, or to change jobs and geographical location. But not all types of mobility are necessarily desirable, and it is easy to think of situations in which there could be 'too much' mobility. Income movements can reflect a substantial degree of riskiness and uncertainty about outcomes. In addition, some large individual reductions may be the result of a willing decision to reduce earnings from employment at a certain age, after accumulating sufficient pension or superannuation benefits. Recognition of the wide range of non-income characteristics, which are usually not available in survey data, makes welfare comparisons much more complex than in the static context.

Within a narrower focus on inequality comparisons, it is not obvious for example whether an increase in inequality is a necessary consequence of 'increased mobility'. Some types of mobility may be equalising while other movements may be disequalising. And a relative income improvement by one person may imply relative downward movements by other individuals, and hence re-ranking. Relevant questions here include: could circumstances arise in which an inequality averse judge would welcome an increase in single-period inequality, when accompanied by a particular type and amount of mobility? An associated question is whether judgements should be made in terms of (some measure of) the sum of incomes over a longer period, rather than being concerned with each individual period and the income changes between periods? The simple sum of incomes over several periods is only one possible measure of longer-term income. An obvious extension is to use a present value, thereby introducing complexities around discounting. And, where discounting is used, it is possible to define an annuity that has the same present value of the actual income stream. Comparisons between income streams having different time profiles of income may differ according to whether a simple sum, a present value or an annuity is used.

The remaining sections therefore move from the earlier discussion of conceptual problems in the measurement of inequality, based on assumptions regarding social welfare functions, to more technical aspects of describing the statistical nature of income mobility. To establish some basic concepts, suppose it is required to obtain a succinct description of the mobility process, and the incomes of a fixed group of individuals in two different periods are available. This gives a two-dimensional joint distribution, characterised by two 'marginal distributions' and two sets of 'conditional distributions'. This is of course the classic 'regression' context: in general for $x$ and $y$ there are two regression lines: these are the relationship between $E\left(y|\, x\right)$ and $x$, and that between $E\left(x|\, y\right)$ and $y$). Of course in the case of joint Normality, these are linear. Just as in the regression context, where there is usually a clear rationale for the choice of dependent and independent variables, in considering mobility the emphasis is usually on the nature of the conditional distributions of second-period incomes, given values of first-period incomes. That is, the perspective is typically 'forward looking'.

It may be expected that some idea of the nature of mobility could be obtained by taking the set of conditional mean second-period incomes, $E\left(x_2|\, x_1\right)$. Based on these, it may be said that there is systematic equalising mobility if the conditional means of initially low-income groups increase by proportionately more than those of the higher-income groups. That is, if the profile of $\left(E\left(x_2|\, x_1\right) - x_1\right)/x_1$ plotted against $x_1$ is downward sloping. Importantly, this profile can be quite different from a profile based on the two marginal distributions. Suppose the percentiles of the marginal distributions are $\left(\varphi_{1,1}, ..., \varphi_{1,100}\right)$ and $\left(\varphi_{2,1}, ..., \varphi_{2,100}\right)$. The profile of $\left(\varphi_{2,i} - \varphi_{1,i}\right)/\varphi_{1,i}$ plotted against $\varphi_{1,i}$ cannot reveal useful information about dynamics: indeed, by regarding individuals as anonymous it completely ignores the way in which the distribution of $x_1$ is transformed into $x_2$. Comparisons of two cross-sectional distributions, even if they contain the same people, cannot give rise to inferences about the nature of mobility.

However, systematically higher percentage income increases for those initially in lower-income percentiles, while suggesting an equalising dynamic process, only tell part of the story. Not everyone in a given conditional distribution experiences the same change as the average. The conditional distributions are characterised by more or less dispersion, arising from a possibly large number of factors affecting income changes. What happens to overall inequality clearly depends on the relative strength of the systematic and non-systematic dynamic components. Before discussing further approaches to describing mobility, the following section explores some of the relationships among the various distributions, using simple models.

# 6 Relationships Among Alternative Distributions

This section explores the relationships among various distributions and the role of income dynamics. This is achieved by focusing on a simplified dynamic process. A basic stochastic process is introduced in Subsection 6.1. Incomes measured over longer accounting periods are discussed in Subsection 6.2. Cross-sectional comparisons of income distributions typically contain multiple overlapping cohorts. The complications raised for inequality comparisons are discussed in Subsection 6.3.

## 6.1 Differential Growth and a Stochastic Component

Consider a process containing both a stochastic component and a systematic component in which changes depend in a straightforward way on the position of individuals relative to the geometric mean. Let $x_{it}$ denote individual $i$'s income in period $t$, and let $\mu_t$ denote the mean of logarithms in period $t$, where $m_t = \exp(\mu_t)$ is the geometric mean. The generating process can be written as:

$$x_{i2} = \left(\frac{x_{i1}}{m_1}\right)^{\beta} \exp(\mu_2 + u_i) \tag{14}$$

where $u_i$ follows a Normal distribution with mean and variance of 0 and $\sigma_u^2$ respectively, so that $u$ is $N(0, \sigma_u^2)$. If $\beta < 1$, there is said to be 'regression towards the (geometric) mean'. Those with $x_{i1} < m_1$ have, on average, larger proportionate increases than those with $x_{i1} > m_1$. Equation (14) can be rewritten as the simple linear regression:

$$(\log x_{i2} - \mu_2) = \beta(\log x_{i1} - \mu_1) + u_i \tag{15}$$

Hence the variance of logarithms of income in period 2, $\sigma_2^2$ , is given by:

$$\sigma_2^2 = \beta^2\sigma_1^2 + \sigma_u^2 \tag{16}$$

An increase in the degree of regression towards the mean (a reduction in $\beta$) reduces $\sigma_2^2$ and an increase in $\sigma_u^2$ increases $\sigma_2^2$. Importantly, a systematic equalising process of regression to the mean ($\beta < 1$) does not necessarily mean that $\sigma_2^2 < \sigma_1^2$. In general, for $\sigma_2^2$ to be less than $\sigma_1^2$, it is necessary for the regression coefficient, $\beta$, to be less than the correlation coefficient between log-incomes in the two periods. Furthermore, the growth rate of each individual's income is:

$$\frac{x_{i2}}{x_{i1}} - 1 = \left(\frac{x_{i1}}{m_1}\right)^{\beta-1} \exp\left(\mu_2 - \mu_1 + u_i\right) - 1 \tag{17}$$

The above kind of dynamic process eventually leads to a stable degree of inequality (as measured by the variance of logarithms of income), $\sigma_S^2$, given by:

$$\sigma_S^2 = \frac{\sigma_u^2}{1-\beta^2} \tag{18}$$

In this model, the use of the variance of log-income as an inequality index is consistent with the basic framework discussed earlier because, for Lognormal distributions, Atkinson's measure is a simple function of this variance.

This process can easily be applied to the simulated distribution used above. Suppose $\beta = 0.8$ and $\sigma_u^2 = 0.2$. Furthermore, suppose that there is underlying growth of $\alpha = 0.1$, that is, 10 per cent, experienced by all individuals. Mean log-income in the second period is $\mu_2 = \mu_1 + \alpha = 3.1$ . From (16) the variance of log-income in the second period is 0.52. Hence, inequality (as measured by this variance) increases despite the systematic regression towards the geometric mean. The question arises of whether this increase in inequality is acceptable to an inequality averse judge, bearing in mind the existence of a systematic equalising component to the changes.

## 6.2 Longer-Period Inequality

The complexity of the process of income dynamics in practice, and the difficulty of obtaining sufficient information about the determinants of individual changes, can suggest an alternative approach which instead concentrates only on inequality measurement. This abstracts from the precise nature of mobility, while producing a more succinct summary. Such an approach begins from a value judgement that longer-period inequality is relevant to a broader judgement regarding inequality. If inequality falls as the accounting period increases, the proportional reduction has sometimes been used as a measure of income mobility.

The simple dynamic process introduced in the previous subsection was used to obtain second-period incomes for each of the 500 simulated individuals. These were used to obtain the sum of incomes measured over two periods: for convenience the second-period incomes were not discounted. The Lorenz curve for the two-period incomes is shown in Figure 5, along with the Lorenz curve for the initial period only. Although $\sigma_2^2 > \sigma_1^2$, the Lorenz curve for the sum of incomes over the two periods lies everywhere closer to the line of equality. Hence, a judge who is concerned only with longer-period inequality is happy to see the increased single-period inequality over time, as part of a

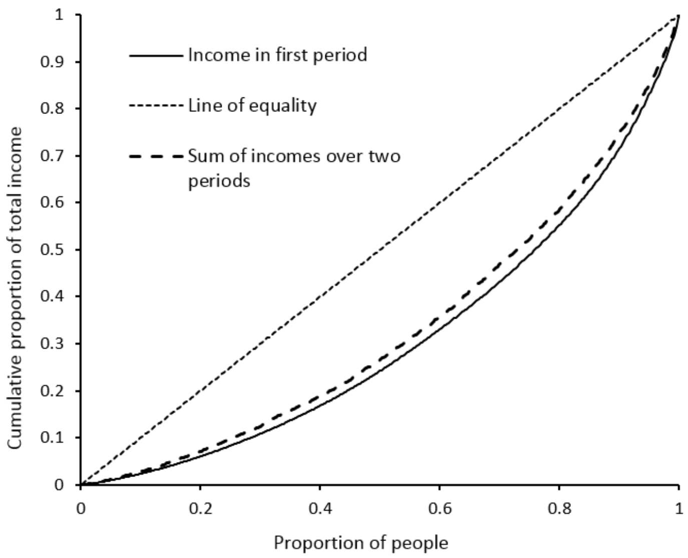


Figure 5: Lorenz Curves for Period 1 and Two-Period Incomes

mobility process that ultimately reduces inequality (relative to the first period) when using a two-year accounting period.

It may be tempting to think that the longer-period inequality necessarily falls as the accounting period increases, and this has often been observed in practice. But comparisons among distributions is far from transparent. Consider the simple case where there is no regression towards the mean and a cohort of individuals, each of whom enters the labour market at the same age. Let $x_{it}$ and $m_t$ denote respectively the income of individual $i$ ($i = 1, ..., n$), and the geometric mean income (defined by $\log m_t = \frac{1}{N} \sum_i \log x_{it}$) in age group $t$ ($t = 1, ..., T$). Define $z_{it}$ as the ratio the logarithm of person $i$'s income to the geometric mean, so that $z_{it} = \log (x_{it}/m_t)$, and for simplicity suppose that, in (14), $\beta = 1$. Hence:

$$z_{it} = z_{i,t-1} + u_{it} \tag{19}$$

Again, $u_{it}$ is a random variable that is independently normally distributed as $N(0, \sigma_u^2)$. In the present context, this simple Markov process is also known as a Gibrat process. Taking variances of (19) gives:

$$\sigma_t^2 = \sigma_1^2 + (t-1)\,\sigma_u^2 \tag{20}$$

The variance of logarithms at age $t$, $\sigma_t^2$, is a linear function of age. Suppose again that all incomes are subject to growth at the constant rate $\alpha$. Arithmetic mean log-income at age $t$, $\mu_t$, is thus:

$$\mu_t = \mu_1 + (t-1)\,\alpha \tag{21}$$

Individual $i$'s total income over $t = 1, ..., T$, say $X_i$, is given (again ignoring discounting) by:

$$X_i = \sum_{t=1}^{T} \exp\left(z_{it} + \mu_t\right) \tag{22}$$

Hence:

$$\log X_i = \left(z_{i1} + \mu_1\right) + \log\left[1 + \sum_{t=2}^{T} \exp\left\{\sum_{s=2}^{t} u_{is} + \alpha\left(t-1\right)\right\}\right] \tag{23}$$

Letting $\Omega$ denote the second term on the right hand side of (23), the variance of logarithms of total income, $\sigma^2_{(T)}$, is equal to:

$$\sigma^2_{(T)} = \sigma^2_1 + V\left[\Omega\right] \tag{24}$$

This depends in a rather awkward way on $\alpha$, $T$ and of course $\sigma^2_u$. Further progress can be made by using the linear approximation for the variance of a nonlinear function of random variables, given by $V\left[f\left(u\right)\right] = f'\left(E\left(u\right)\right)^2 V\left(u\right).$ Noting that all $u_{it}$ are from the same distribution with mean 0, an approximation to (24) is given by:

$$\sigma^2_{(T)} = \sigma^2_1 + \sigma^2_u \left[\frac{\sum_{t=2}^{T} \exp\left\{\alpha\left(t-1\right)\right\}}{1 + \sum_{t=2}^{T} \exp\left\{\alpha\left(t-1\right)\right\}}\right]^2 \tag{25}$$

As expected, if there is no mobility, $\sigma^2_u = 0$ and $\sigma^2_{(T)}$ is the same as that in the first year. In general, it is possible for, say, the ranking of groups according to $\sigma^2_{(T)}$ to change as the length of time over which incomes are measured, $T$, is gradually increased.

## 6.3 Multiple Cohorts and Cross-Sections

The cross-sectional distribution consists of individuals from many cohorts, at different stages of the life cycle. For simplicity, suppose that each cohort has the same length of life, $T$, and has similar values of $\alpha$, $\sigma^2_1$ and $\sigma^2_u$. Aggregating over many cohorts requires an age distribution to be specified. Suppose that $h_t$ denotes the proportion of the population who are aged $t$ in a given time period, so that $\sum_{t=1}^{T} h_t = 1$. The mean of logarithms in the cross section, $\mu_X$, is therefore equal to $\sum_{t=1}^{T} \mu_t h_t$ and the variance of logarithms, $\sigma^2_X$, is given, following the standard decomposition into within-age and between-age components, by:

$$\sigma^2_X = \sum_{t=1}^{T} h_t \sigma^2_t + \sum_{t=1}^{T} h_t \left(\mu_t - \mu\right)^2 \tag{26}$$

For $t \geq 2$, the terms in $\mu_t$ and $\sigma^2_t$ are given by simple expressions involving $\alpha$ and $\sigma^2_u$, so it is possible to expand equation (26). For example, expanding the first term, this becomes:

$$\sigma_X^2 = \sigma_1^2 + \sigma_u^2 \sum_{t=2}^{T} (t-1)\, h_t + \sum_{t=1}^{T} h_t \left(\mu_t - \mu\right)^2 \tag{27}$$

Hence, $\sigma_1^2$ provides a lower bound to the variance of logarithms in both cross-sectional and lifetime contexts, and $\sigma_u^2$ affects the variance of logarithms in the cross-sectional distribution only through the second term in (27).

The analytics rapidly become much more complex if the various simplifying assumptions are relaxed, but the results can be used to examine whether the following two hypotheses are likely to hold. First, if two countries have broadly similar degrees of income mobility, cross-sectional comparisons can provide a good indication of differences in lifetime inequality. Second, if an increase in cross-sectional inequality in a country is associated with an increase in income mobility, lifetime inequality increases by less than cross-sectional inequality. Both these hypotheses may at first sight appear to be reasonable, and if true they would be very useful. The first extends the range of international comparisons that can be made with limited data, and the second suggests, for example, that greater labour market flexibility has a smaller impact on inequality when a longer period measure is considered.

A comparison of equations (25) and (26) shows immediately that in the extreme case where there is no relative income mobility within cohorts, $\sigma_u^2 = 0$ and $\sigma_{(T)}^2 = \sigma_1^2$, while $\sigma_X^2 > \sigma_1^2$ because cross-sectional inequality depends on the steepness of the age-income profile and the age distribution. Two countries can therefore both have no mobility, but the cross-sectional distributions can give quite misleading indications of lifetime inequality, depending on the values of $\alpha$ and the distributions of $h_t$. Similarly, for common non-zero values of $\sigma_u^2$, the cross-section can be equally misleading. This argument therefore shows that the first hypothesis above is rejected.

To consider the second hypothesis, differentiate both (25) and (27) with respect to $\sigma_u^2$. It can be seen that $\partial\sigma_X^2/\partial\sigma_u^2$ depends only on the form of the age distribution, while the term $\partial\sigma_{(T)}^2/\partial\sigma_u^2$ depends only on the parameter $\alpha$ (and of course $T$). Hence there is no reason why an increase in mobility, reflected in an increase in $\sigma_u^2$, should be expected to increase lifetime inequality by less than cross-sectional inequality because they depend on quite different factors. These examples show that even in a highly simplified model, great care is needed in making comparison on the basis of cross-sectional distributions. There is no easy substitute for obtaining genuine longitudinal data. Having examined the relationships among various distributions, for simple dynamic processes, the following section turns to consideration of how mobility, viewed in terms of differential income growth, can be illustrated diagrammatically.

# 7 Illustrating Income Mobility

It has been seen that the valuable Lorenz curve, based on cumulative values, provides a smooth curve and convenient picture of the nature of the dispersion of incomes. This is more transparent than the basic histogram, whose appearance can depend on the somewhat arbitrary choice of income classes used. The question arises of whether a corresponding curve can be defined in the mobility context, to illustrate systematic equalising or disequalising income changes. First, as with the single-period distribution, it is clear that a simple scatter diagram is not very informative. For example, using the simulated data described above, the resulting basic scatter diagram is shown in Figure 6. For larger populations this would clearly consist of a large dark area, where individual observations cannot be distinguished.

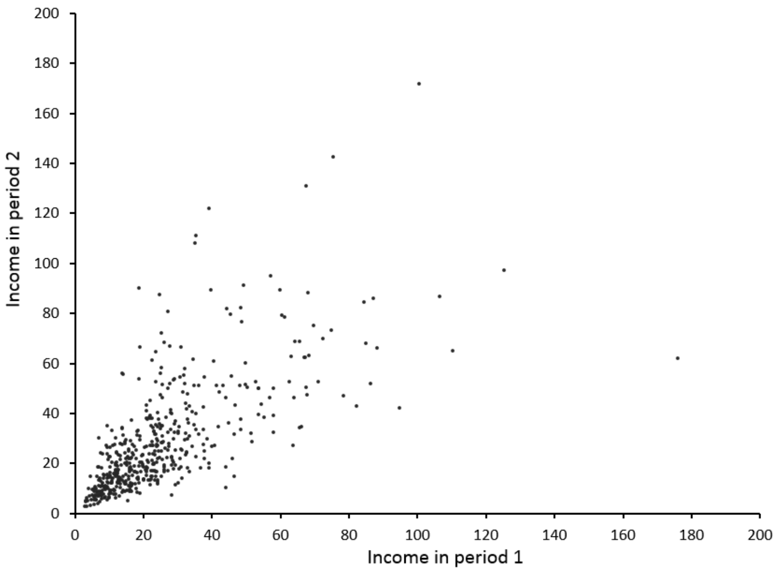


Figure 6: Scatter Plot of Simulated Incomes in Two Periods

Given empirical data on incomes of a constant group of people in two different years, one popular method of examining the nature of the dynamic process is to produce a transition matrix, describing the movement from particular ranges of the initial distribution to ranges of the distribution in period 2. To illustrate the nature of a transition matrix, the dynamic process described above was applied to the 500 initial simulated incomes, and used to obtain the proportions moving among different income classes. For simplicity, consider only four groups defined using the lower and upper quartiles, $q_U$ and $q_L$ and the median, $q_M$ in each period. The transition matrix is given in Table 1.

The income movements are from rows to columns of the table. For example, the

Table 1: Transition Matrix

| Period 1 | Period 2 income class | | | | Total |
|---|---|---|---|---|---|
| income class | $0 - q_L$ | $q_L - q_M$ | $q_M - q_U$ | $> q_U$ | |
| $0 - q_L$ | 0.61 | 0.26 | 0.13 | 0.00 | 1.00 |
| $q_L - q_M$ | 0.30 | 0.42 | 0.24 | 0.05 | 1.00 |
| $q_M - q_U$ | 0.08 | 0.27 | 0.37 | 0.28 | 1.00 |
| $> q_U$ | 0.01 | 0.10 | 0.23 | 0.66 | 1.00 |
| Total | 1.00 | 1.00 | 1.00 | 1.00 | |

value of 0.26 in the first row indicates that, of those who started in the lowest quartile of the distribution, 26 per cent moved up to an income level between the second period's lower quartile and median. Row and column sums are thus all unity. In some contexts it may be desirable to add a column for 'exits' and a row for 'entrants' to the population group considered. Where deciles are used, more detail is obviously revealed, although the nature of mobility may become less transparent. Hence, further indices, such as the proportion of upward and downward movements, may be obtained. However, there is clearly a need for a more succinct measure of mobility, and way of visualising the nature of movements 'at a glance'.

Section 5 discussed the joint distribution, and the possible nature of a profile of the proportional change, $\left(E\left(x_2 | x_1\right) - x_1\right)/x_1$, plotted against $x_1$. A downward sloping profile is suggestive of systematic equalising mobility. However, a richer description of mobility can be obtained using a somewhat different treatment of proportional income changes, as shown in the following section.

# 8 The 'Three "I"s of Mobility'

This section describes a curve, called the 'Three "I"s of Mobility', or TIM, curve for illustrating several characteristics of income mobility. Define the logarithm of income, $y_i = \log x_i$, for individuals $i = 1, ..., n$. Hence $y_{i,t} - y_{i,t-1}$ is (approximately) person $i$'s proportional change in income from period $t-1$ to $t$. With log incomes ranked in ascending order, plot $\frac{1}{n}\sum_{i=1}^{k}\left(y_{i,t} - y_{i,t-1}\right)$ against $h = \frac{k}{n}$, for $k = 1, ..., n$. Thus the TIM curve plots the cumulative proportional income change per capita against the corresponding proportion of individuals, $h$.

A TIM curve allows focus on the mobility of a particular group of low-income individuals: those with incomes below $x(h)$, for the proportion, $h$, of the population. In this framework $h$ captures the incidence of the particular group of concern. Similarly, the intensity and inequality dimensions of mobility in terms of income growth are reflected in the shape of the TIM curve, by analogy with the TIP curve.

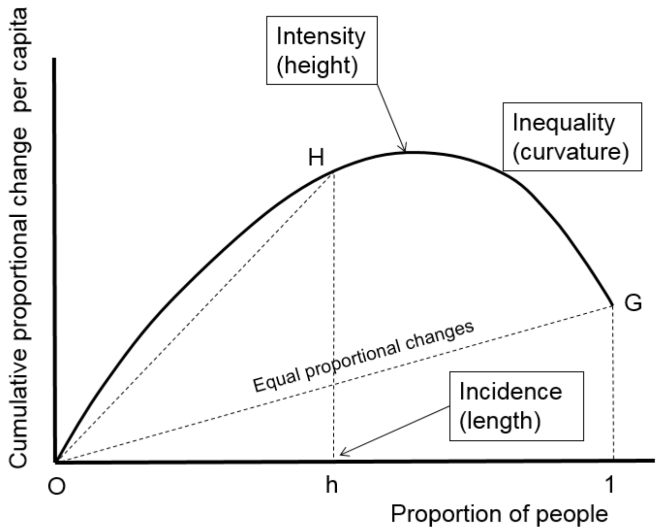


Figure 7: A TIM Curve

A hypothetical example of a TIM curve is shown in Figure 7, with $h = k/n$ on the horizontal axis. This reflects a situation in which relatively lower-income individuals receive proportional income increases which are greater than that of average (geometric mean) income. Hence the TIM curve, OHG, lies wholly above the straight line OG. The TIM curve in practice is not necessarily smooth and concave over the whole range, like the hypothetical one shown in Figure 7. There can be income reductions as well as increases and, depending on how they are spread across the initial income distribution, and the curve can be more or less 'jagged' depending on the extent of non-systematic income-related changes.

If all incomes increase by the same proportion, the TIM curve is the straight line OG. The height, G, indicates the average growth rate of the population as a whole, with the height, H, indicating the average growth rate for those below $x(h)$. Furthermore, inequality is reflected in the degree of curvature. For example, the curvature of the arc OH relative to the straight line OH indicates that lower income individuals have higher (more unequal) growth than those individuals to the left of, but closer to, $h$.

Suppose interest is focussed on those below the $h^{th}$ percentile, indicated in Figure 7. There is less inequality of mobility within the group below $h$, shown by the fact that the TIM curve from O to H is closer to a straight line than the complete curve OHG. The TIM curve also shows that the income growth of those below $h$ is larger than that of the population as a whole. The average growth rate among the poor (the intensity of their growth) is given by the height H.

For comparison purposes, a normalisation of the TIM curves is required. For example, comparing the income mobility experienced across different periods, the arithmetic mean income growth rate, $g$, is likely to vary across periods, such that the height of point G in Figure 7 differs. This can make comparisons of the degree of inequality of mobility across periods difficult. In this case equivalent normalised TIM curves, or nTIM curves, can be obtained where each TIM is normalised by the average growth rate for each period. With normalisation, $M_{h,t}$ reaches a value of 1 at $h = 1$, though $M_{h,t}$ values can exceed 1 at lower values of $h$, as illustrated in Figure 7. This normalisation allows the degree of concavity or convexity of each TIM curve to be directly compared.

Using the simulated data for 500 individuals in two periods, described above, the normalised TIM curve is shown in Figure 8. This curve lies above the straight line (of equal proportional income changes) over its whole length, reflecting the assumed regression towards the geometric mean ($\beta = 0.8$). Yet a certain amount of jaggedness is apparent in view of the high degree of random variation ($\sigma_u^2 = 0.2$).

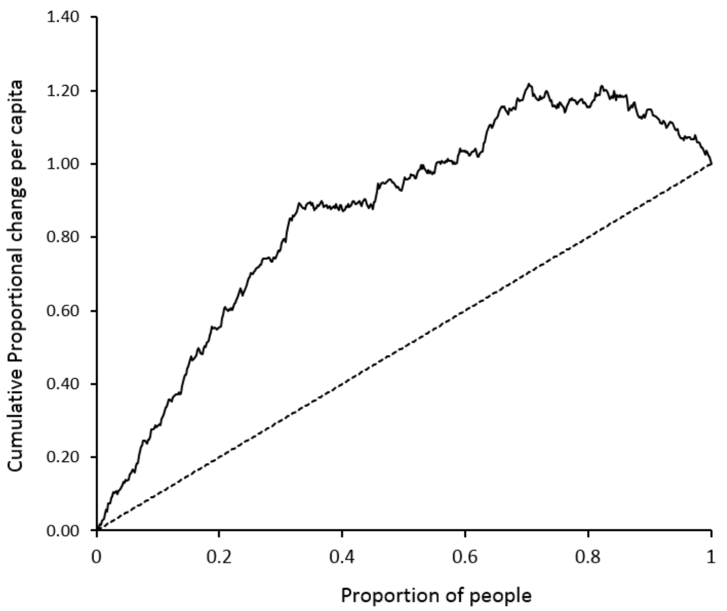


Figure 8: The Normalised TIM Curve for Simulated Data

When discussing the measurement of inequality above, stress was placed on the approach whereby a measure reflects the explicitly assumed distributional value judgements of an independent disinterested judge. In practice, results are reported for a limited range of assumptions, and a very limited approach is taken to the question of what might be considered relevant non-income characteristics of individuals: complex theoretical difficulties arise when individuals are heterogeneous. When considering relative income changes taking place over time – which can be subject to factors involving age and calender time, and which may in principle vary among cohorts – it becomes even harder to specify the nature of value judgements that may be applied in evaluating a dynamic process. This is because, as discussed above, there are many factors, endogenous and exogenous, which could give rise to mobility. Some changes can be due to misfortune (illness, job loss, adverse market conditions), while others can arise from fortunate exogenous shocks, and yet others can arise from the kind of changes discussed above. Faced with this difficulty, it is nevertheless of value to attempt to determine whether systematic forces have been at work, and whether these differ between countries or over time. A possible approach is discussed in the following subsection.

# 9 The TIM Curve and Equalising Mobility

In the case of the Lorenz curve, it has been seen that the definitions of complete equality and inequality (whereby only one person has a positive income while all others are zero) are straightforward to define and envisage. In the context of mobility as differential growth, attempting to define the extremes is not so straightforward. However, one extreme case, that of a relative inequality-preserving mobility process, is simple.

Consider one extreme form of inequality-reducing mobility, where income changes are either zero or positive. If only the poorest person has an income increase, while all other incomes remain unchanged, the normalised TIM curve is simply a horizontal line, after following the vertical axis up to 1. But this is an arbitrary case. Another possibility is the set of proportional income changes which produce equal incomes in the second period, equal to the actual average in that period. In terms of incomes, $x_{2,i}$ and $x_{1,i}$, for person $i = 1, ..., n$ in periods 2 and 1, for growth rates, $g_i$, and second-period arithmetic mean income, $\bar{x}_2$, this requires $x_{2,i} = x_{1,i}(1 + g_i) = \bar{x}_2$, or (for strictly positive initial incomes):[2]

$$g_i = \left(\frac{\bar{x}_2}{x_{1,i}}\right) - 1 \tag{28}$$

However, these $g_i$s produce a TIM curve having an average growth rate that differs from the actual average growth rate, which means that the normalised version adjusts the cumulative growth rates per capita by a different amount from that used to obtain the actual TIM curve. Suppose it is required to have equal incomes, $\widetilde{x}$, in the second period *and* an average growth rate equal to the actual rate, $g$. Using logarithmic changes to approximate proportional changes:

$$g = \frac{1}{n}\sum_{i=1}^{n}(\log x_{2,i} - \log x_{1,i}) \tag{29}$$

[2]If changes are expressed as log-income-changes, then $g_i = \log \bar{x} - \log x_{1,i}$. This uses the approximation $\log(1 + g_i) = g_i$.

The $\widetilde{x}$ and associated $g_i$ must now satisfy:

$$g_i = \log \widetilde{x} - \log x_{1,i} \tag{30}$$

and:

$$\frac{1}{n}\sum_{i=1}^{n} g_i = g \tag{31}$$

Substituting (30) into (31) and equating with (29) gives:

$$\log \widetilde{x} = \frac{1}{n}\sum_{i=1}^{n} \log x_{2,i} \tag{32}$$

Furthermore, substitution into (30) gives the set of growth rates needed:

$$g_i = \left(\frac{1}{n}\sum_{i=1}^{n} \log x_{2,i}\right) - \log x_{1,i} \tag{33}$$

Figure 9 illustrates two normalised TIM curves for a given initial income distribution in period 1. The solid line is the actual nTIM curve, and the higher dashed line is the hypothetical curve which would arise from the application of proportional income changes according to equation (33).

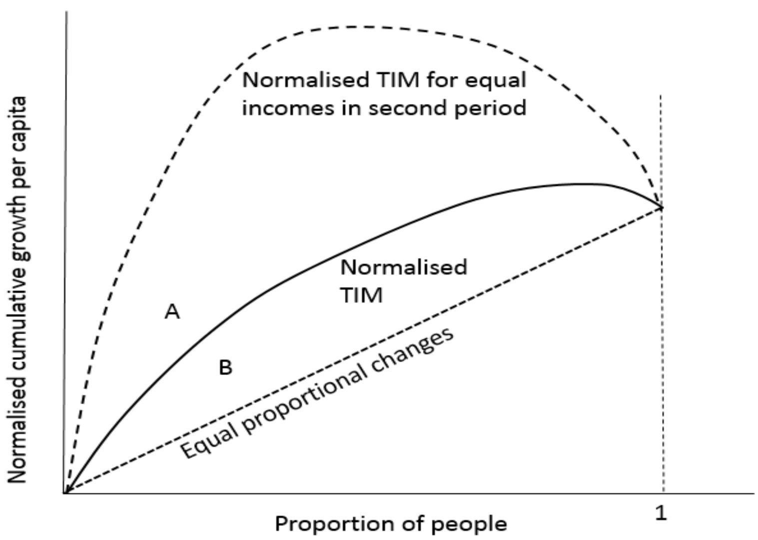


Figure 9: Actual nTIM Curve and nTIM corresponding to Equal Second-Period Incomes and an Average Growth Rate Equal to the Actual Rate

The question is whether a useful measure of the degree of systematic equalising mobility can be obtained in this diagram. Clearly, the area $B$ (between the nTIM curve and the diagonal line of equal proportional changes) alone does not provide an appropriate measure, since the scope for equalising differential income growth depends on the initial income distribution. The same dynamic process (in terms, say, of the relationship between the proportional growth rate and initial income) gives a smaller area for a relatively more equal distribution than for a more unequal distribution (for the same overall income growth rate). The maximum area is $A + B$, where $A$ is the area between the actual nTIM and the hypothetical 'fully equalising-mobility' nTIM. Consider a measure of the degree of equalising mobility, $M_E$, defined as:

$$M_E = \frac{B}{A + B} \tag{34}$$

The maximum value this can take is 1 while the minimum is 0 (when relative incomes do not change). The example shown is one in which there is systematic second-period-equalising mobility, in that the nTIM curve is everywhere above the diagonal nTIM of equal proportional changes, which is the dominant case in empirical applications.

It is important to recognise that the equalising TIM or nTIM curve is not uniquely defined, as it depends on the form of the initial distribution. A more equal distribution in period 1 produces a lower and less-concave curve. Nevertheless, the fact that a measure of equalising mobility depends on the initial distribution is not really surprising. For example, it corresponds to the fact that, in a slightly different context (though one involving movement from one distribution to another), income tax progressivity measures depend on the initial or pre-tax income distribution.

# 10 Poverty Persistence

This section turns to the question of how poverty persistence can be shown diagrammatically, reflecting the extent to which upward income mobility between two periods shifts individuals from below, to above, a given income poverty threshold, say $x_p$. Since the TIM curve illustrates the *cumulative* extent of mobility for those below the associated $h_p$, it can show whether the incomes of an initially poor group *on average* grew sufficiently to escape from poverty, but it cannot directly illustrate the extent to which *individuals* below $h_p$ escape from poverty.

Some individuals within this group in time period, $t - 1$, may experience sufficient income growth between $t - 1$ and $t$ to raise their income levels above $x_p$. Assume further that the poverty income threshold is constant in both years. Let $g_i = \frac{dx_i}{x_{i,t-1}}$ denote individual $i$'s proportional income growth between $t - 1$ and $t$. The condition required for those individuals for whom $x_{i,t-1} < x_p$ to move out of poverty is given by:

$$g_i > \frac{x_p}{x_{i,t-1}} - 1 \tag{35}$$

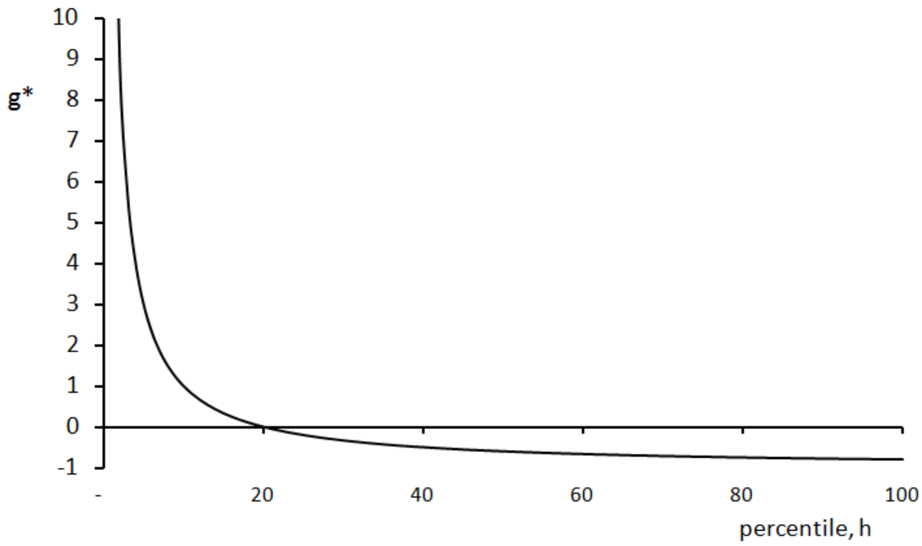


Figure 10: Income Growth Rates Required to Escape and Avoid Entering Poverty

More generally, all individuals can be allocated to one of four groups based on their values of $g_i$ and $x_{i,t-1}$, as shown in Table 2.

Table 2: Poverty Persistence

| | In poverty | Out of poverty |
|---|---|---|
| Move | $g_i < \frac{x_p}{x_{i,t-1}} - 1$<br>$x_{i,t-1} > x_p$ | $g_i > \frac{x_p}{x_{i,t-1}} - 1$<br>$x_{i,t-1} < x_p$ |
| Persist | $g_i < \frac{x_p}{x_{i,t-1}} - 1$<br>$x_{i,t-1} < x_p$ | $g_i > \frac{x_p}{x_{i,t-1}} - 1$<br>$x_{i,t-1} > x_p$ |

The groups are separated by values of $g_i$ and $x_{i,t-1}$, denoted $g_i^*$ and $x_{i,t-1}^*$ respectively, given by $g_i^* = \frac{x_p}{x_{i,t-1}} - 1$ and $x_{i,t-1}^* = x_p$. These can be illustrated by a variant of the TIM curve. As discussed above, the TIM curve plots, for incomes in ascending order, cumulative proportional income changes per capita against the corresponding proportion of people, $h$. Consider an alternative diagram in which individual income growth rates, $g_i$ and $g_i^*$ (which depends on the proportional difference between $x_{i,t-1}$ and $x_p$) are plotted against $h$, for any given income poverty threshold, $x_p$, and associated $h_p$.

Figure 10 plots the values of $g^*$ against $h$, for a poverty income threshold, $x_p$, such that $h_p = 0.2$: that is, 20 per cent of the population are below the poverty line. Hence, the $g^*$ profile crosses the $x$-axis at $h = 0.2$. This is referred to as a *poverty persistence curve*, $g^*(h_p)$, which is defined for a given value of $h_p$.

To the left of $h_p = 0.2$, growth rates greater than $g^*$ from period $t-1$ to period $t$ are sufficient to move individuals out of poverty, that is, to an income level which places them above $h_p$ in the population in time period, $t$. Conversely, for those who are placed to the right of $h_p = 0.2$ in period $t-1$, growth rates less than $g^*$ are sufficiently negative to move the individual into poverty in period $t$. The profile of critical values, $g^*$, approaches an asymptote of $-1.0$ (or $-100$ per cent); that is, as incomes become very large relative to $x_p$, the required (negative) growth rate to move such individuals into poverty from period $t-1$ to period $t$ approaches $-100$ per cent.

# 11 Positional Changes

The discussion above has been of income mobility defined in terms of relative income growth. The extent to which individuals change their rank order in the distribution – when incomes are arranged in ascending order – may also be regarded as an important dimension of mobility. The question arises of whether a diagram similar to the TIM curve can be helpful in this context. Individuals can move to higher or lower rank positions, so that explicit treatment of the *direction* of change becomes important. Defining a re-ranking mobility index, it is first necessary to decide whether negative re-ranking (dropping down the ranking) is treated symmetrically with positive (upward) movement within the ranking. A second issue concerns the choice of whose mobility is to be included.

In the following discussion, individuals are ranked in ascending order of initial incomes, $x_{i,0}$, so that ranks $i = 1, ..., n$ are for individuals from the lowest to the highest income. The initial period is denoted 0, and initial ranks may be defined as $R_{i,0} = i$. Consider, as in previous sections, the case where it is desired to measure the extent of mobility of a subset of individuals, $k \leq n$, with the lowest initial incomes, and let $\Delta R_i = R_{i,1} - R_{i,0} = R_{i,1} - i$ denote the change in the rank order of the person who initially has rank, $i$. Three treatments of re-ranking are possible, all related to how negative, or downward, re-ranking is treated. Firstly, negative re-ranking could be treated symmetrically with positive re-ranking such that positional mobility is defined in *net* terms, that is, positive changes in rank net of any negative changes for those with $i \leq k$. This is referred to as '*net* re-ranking'. Secondly, negative movement in the ranking could be ignored, which simply involves setting $\Delta R_i = 0$ when $\Delta R_i < 0$. This is referred to as '*positive* re-ranking'. Thirdly, re-ranking may be measured in *absolute* terms in which all re-ranking is measured as a positive value. This is referred to as '*absolute* re-ranking'.

The choice among these three measures depends on the question of interest. For example, if interest is focussed on those below the poverty line as a group, then it may be desirable to balance any upward mobility by some of those in poverty with downward mobility of others in poverty, in order to gain an indication of the overall experience of the group. This suggests a focus on net mobility in this case. If movement *per se* is the mobility concept of interest, then a non-directional measure such as absolute re-ranking is more relevant. Positive re-ranking quantifies only those who are moving up, a common metric when assessing the persistence of low income or poverty status for a sub-set of individuals or households.

The three re-ranking indices for individual, $i$, are defined formally as follows:

$$M_i^{net} = \Delta R_i, \tag{36}$$

$$M_i^{pos} = \Delta R_i\big|_{\Delta R_i \geqslant 0}, \tag{37}$$

$$M_i^{abs} = \left|\Delta R_i\right|. \tag{38}$$

Cumulated across the $k$ lowest-income individuals in period 0, the corresponding aggregate re-ranking indices are given by:

$$M^{net}(k) = \sum_{i=1}^{k} M_i^{net} = \sum_{i=1}^{k} (R_{i,1} - R_{i,0}), \tag{39}$$

$$M^{pos}(k) = \sum_{i=1}^{k} M_i^{pos} = \sum_{i=1}^{k} (R_{i,1} - R_{i,0}) \quad \text{for} \quad \Delta R_i \geqslant 0, \tag{40}$$

and:

$$M^{abs}(k) = \sum_{i=1}^{k} M_i^{abs} = \sum_{i=1}^{k} \left|R_{i,1} - R_{i,0}\right|. \tag{41}$$

The absolute re-ranking case may be thought of as describing the extent of overall positional change within the relevant range of the income distribution.

To examine the 'three Is' properties similar to the TIM, but based on the indices in (39), (40) and (41), one approach would be to plot the value of the relevant $M(k)$ index against the cumulative fraction of the population, $h = k/n$. However, a difficulty with the indices in (39) to (41) is that they are not scale independent, since they depend on the population size: more re-ranking is possible in larger populations. One solution would be to scale the three $M(k)$ indices by $n$. However, re-scaling, by $(n/2)^2$, yields normalized values, $m(k)$, that lie between zero and one (or zero and two for absolute re-ranking). These may then readily be plotted against $0 \leqslant h \leqslant 1$. Hence it is possible to define a cumulative re-ranking curve, similar to the TIM curve for relative income mobility, that plots $m(k)$ against $h$. Furthermore, an individual's opportunity for re-ranking is partly determined by their initial position. Those among the lowest ranks have less opportunity to move down, other things equal, than those higher up, and *vice versa*. It is therefore useful to consider the maximum re-ranking possible for each individual. Actual re-ranking may then be compared with these maximum values for any given $h$.

On maximum re-ranking, consider a population of $n$ individuals, each with a different income level. In period 0, they are ranked, $R_{i,0} = i$, for $i = 1, \ldots, n$, representing the lowest to the highest incomes. Two polar cases are the maximum and minimum degrees of mobility possible. The former is defined here as a complete ranking reversal, $\Delta R_i(\max)$, which implies:

$$M_i(\max) = \Delta R_i(\max) = R_{i,1}(\max) - R_{i,0} = n + 1 - 2R_{i,0}. \tag{42}$$

For large $n$, this is approximated by $n - 2R_{i,0}$. Where it is desired to measure the extent of re-ranking of the subset of individuals, $k \leq n$, with the lowest incomes, the cumulative maximum re-ranking index for the net mobility case , $M^{net}(\max, k)$, is given by:

$$M^{net}(\max, k) = \sum_{i=1}^{k} M_i^{net}(\max) = \sum_{i=1}^{k}(n + 1 - 2R_{i,0}). \tag{43}$$

Using the sum of an arithmetic progression, whereby $\sum_{i=1}^{k} R_{i,0} = 1 + 2+ \ldots + k = k(k+1)/2$, equation (43) becomes:

$$\begin{aligned} M^{net}(\max, k) &= \sum_{i=1}^{k}(n + 1 - 2R_{i,0}) = k(n+1) - k(k+1) \\ &= k(n-k). \end{aligned} \tag{44}$$

Hence, for example, if $n = 100$, each integer, $i = 1, ..., n$, represents a percentile of the distribution. If interest focuses only on the poorest individual (so that $k = 1$), maximum net re-ranking is given by $M^{net}(\max, 1) = (100 - 1) = 99$; when $k = 2$, $M^{net}(\max, 2) = 2(100 - 2) = 196$; and so on. More generally, since maximum re-ranking (complete ranking reversal) involves all those below the median individual changing positions with those above the median, it follows from (44) that the maximum value of $M^{net}(\max, k)$ as $k$ increases is obtained for $k = n/2$, yielding $M^{net}(\max, n/2) = (n/2)^2$, for large $n$.

This measure serves to highlight the scale dependence of both $M^{net}(k)$ and $M^{net}(\max, k)$: larger populations imply larger values of both indices. These could be normalized to create a form of per capita index by dividing by $n^2$ such that, from (44), the index would become: $M^{net}(\max, k)/n^2 = h(1 - h)$. The maximum value would be reached at $h = 0.5$, where the index equals 0.25. However, to yield an index with a maximum value of 1 at $k = n/2$, it is preferable to divide by $(n/2)^2$. That is,

$$m^{net}(\max, k) = 4M^{net}(\max, k)/n^2, \tag{45}$$

where lower-case $m$ denotes this normalized measure. Using (44) and remembering that $k = nh$:

$$m^{net}(\max, nh) = 4h(1 - h). \tag{46}$$

A similar exercise for positive re-ranking, $M^{pos}(\max, k)$, shows that $M^{pos}(\max, k)$ also reaches a maximum as $k$ increases, of $M^{pos}(\max, n/2) = n^2/4$, since all individuals below $n/2$ experience positive re-ranking in this maximum case. However, above $k = n/2$, as more above-median individuals are included within $k$, their re-rankings are now given by $\Delta R_i = 0$; hence, the cumulative index, $M^{pos}(\max, k)$, remains unchanged for $k > n/2$. Thus a similarly re-scaled $m^{pos}(\max, k)$ may be defined analogously to (45) to yield a positive re-ranking index where $0 \leqslant m^{pos}(\max, k) \leqslant 1$.

Finally, for the absolute re-ranking case in (41), $M^{abs}(\max, k)$, it can be shown that, as with the other cases, this increases as $k$ increases from $k = 1$ to $k = n/2$ to reach $M^{abs}(\max, n/2) = (n/2)^2$. However, this represents a point of inflection rather than a maximum, since inclusion of the absolute value of above-median individuals' re-ranking in $M^{abs}(\max, k)$, ensures that $M^{abs}(\max, k)$ continues to increase for $k > n/2$, reaching $M^{abs}(\max, n) = n^2/2$ at $k = n$. As a result, an absolute re-ranking index $m^{abs}(\max, k)$ obtained by re-scaling by $(n/2)^2$ lies between zero and two.

Finally, to compare actual and maximum re-ranking mobility, the expressions for actual mobility in (39) to (41) can be similarly re-scaled or normalized by $(n/2)^2$ to obtain actual aggregate re-ranking mobility expressions, $m^{net}$, $m^{pos}$, and $m^{abs}$, given in each case by:

$$m = 4M/n^2. \tag{47}$$

Thus, $0 \leqslant m^{net}$, $m^{pos} \leqslant 1$ and $0 \leqslant m^{abs} \leqslant 2$. This suggests a convenient illustrative device for positional mobility, a cumulative re-ranking curve, similar to the TIM curve for relative income mobility, that plots alternative $m$s against $h$. This is explored in the next section using income data for three large longitudinal samples of New Zealand individual taxpayers over 1998 to 2010. First, the next subsection shows $m(\max, k)$ profiles and introduces an alternative illustration based on the ratio of actual to maximum re-ranking: the re-ranking ratio, $RRR$.

## 11.1 Maximum Re-Ranking profiles

Profiles for the three (re-scaled) maximum re-ranking cases discussed above, $m^{net}(\max, k)$, $m^{pos}(\max, k)$, and $m^{abs}(\max, k)$, are plotted against $h = k/n$ in Figure 11. This shows the distinct non-linear shape of the maximum profiles, whichever definition of positional mobility is adopted. As expected, the net re-ranking profile displays a parabolic shape which differentiation of (46) reveals has a slope of $4(1-2h)$, that equals zero at $h = 0.5$ (the $50^{th}$ percentile), thereafter declining symmetrically to a slope of $-4$ at $h = 1$. The equivalent positive re-ranking profile also reaches a maximum at the $50^{th}$ percentile but remains constant thereafter, while the absolute re-ranking profile displays a sigmoid shape, reaching a local point of inflection where $m^{abs}(\max, 0.5n) = 1$ at the $50^{th}$ percentile, but then rising at an increasing rate till $m^{abs}(\max, n) = 2$ at $h = 1$.

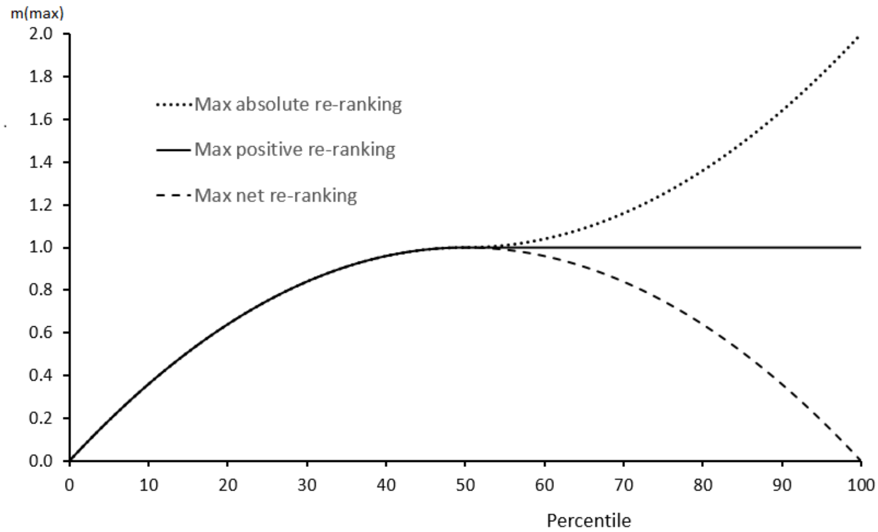


Figure 11: Maximum Re-ranking Profiles

The maximum re-ranking indices in Figure 11 are invariant to population size, but they vary with the population percentile of interest, $h$. Thus, the *scope* or opportunity for a given degree of actual re-ranking clearly also varies with $h$. A natural index of interest therefore is the ratio of actual to maximum mobility at each percentile, $h$. This is referred to as the re-ranking ratio, $RRR$, and can be calculated for net, positive and absolute re-ranking. For example, the net re-ranking case is given by:

$$RRR^{net} = \frac{m^{net}}{m^{net}(\max)} = \frac{M^{net}}{M^{net}(\max)}, \tag{48}$$

where the numerator and denominator are given respectively by (47) and (45), or by (39) and (44). This ratio can also be plotted against $h$ to identify how the extent of mobility changes by cumulative percentile of the population relative to the maximum possible *for that percentile*.

Recognising these differences in maximum re-ranking is important when interpreting differences in actual re-ranking for different values of $h$. In particular, a smaller value of $m^{net}$ at $h = 0.1$, compared to $m^{net}$ at $h = 0.3$, for example, may be partly or entirely due the fact that individuals up to $h = 0.1$ cannot achieve the higher $m^{net}$ observed at $h = 0.3$, even in the absence of other constraints on re-ranking mobility.

# 12 Conclusions

This paper has explored some basics of income inequality measurement, and examines the challenges involved when extending the accounting period to more than a single year. When discussing inequality measurement, the emphasis was on conceptual aspects, in particular the way in which the implications of adopting alternative explicit value judgements can be examined. But in considering incomes over several periods, a greater emphasis has been given to more technical descriptive methods of analysis.

It may be said that the discuss has concentrated on two contexts where a fundamental value judgement, that of anonymity (or 'symmetry'), needs to be rejected. Anonymity requires that the non-income characteristics and identities of individual income earners are not considered to be relevant. However, and despite the fact that the 'individual' is in some sense a fundamental unit of analysis, many individuals belong to families or households, within which some form of income sharing takes place. In comparing the incomes of two individuals, it therefore becomes important to know the family or household to which those individuals belong: they cannot be treated as anonymous. Second, in a multi-period context in which incomes fluctuate over time, it is necessary to be able to trace the precise income paths followed by different individuals, who may also move between family or household types. Individual income profiles rely on being able to identify those individuals: they cannot be anonymous.

It is no accident that the study of inequality is a topic within the aegis of economics, rather than statistics, although some knowledge of basic statistical concepts is required. Specifically, the analysis of inequality falls within the realm of welfare economics. Crucially, this paper has stressed that income comparisons are not value free: the economic approach involves investigating the implications of adopting explicitly stated distributional value judgements. The readers, or users of the analysis, can decide if those values correspond to anything like their own. If not, at least the results give an idea of why their policy prescriptions may differ. There is thus much value in investigating results for as wide a range of assumptions as possible, so that the sensitivity of results to different value judgements becomes clear.

# References

[1] Creedy, J. (1975) Aggregation and the distribution of income. *Oxford Bulletin of Economics and Statistics,* 37, pp. 91-101.

[2] Creedy, J. (1977) The distribution of lifetime earnings. *Oxford Economic Papers*, 29, pp. 412-429. [Reprinted in *Income Distribution: Description, Measurement, Shape, Dynamics,* (ed. by. M. Sattinger), Vol. I, pp. 468-485). Cheltenham: Edward Elgar (2001)]

[3] Creedy, J. (1985) *Dynamics of Income Distribution.* Oxford: Basil Blackwell, pp. xiii+170. [Reprinted in 1993, Aldershot: Gregg Revivals]

[4] Creedy, J. (1997a) Inequality, mobility and income distribution comparisons. *Fiscal Studies,* 18, pp. 293-302.

[5] Creedy, J. (1997b) Lifetime inequality and tax progressivity with alternative income concepts. *Review of Income and Wealth*, 43, pp. 283-295.

[6] Creedy, J. (2016) Interpreting inequality measures and changes in inequality. *New Zealand Economic Papers*, 50, pp. 177-192.

[7] Creedy, J. (2017) Alternative distributions for inequality comparisons. *Australian Economic Review*, 50, pp. 484-497.

[8] Creedy, J. (2023a) Comparing income distributions using Atkinson's measure of inequality. *Australian Economic Review* (forthcoming).

[9] Creedy, J. (2023b) Distributional comparisons using the Gini inequality measure. *Australian Economic Review* (forthcoming).

[10] Creedy, J. and Gemmell, N. (2018) Income dynamics, pro-poor mobility and poverty persistence curves. *Economic Record*, 94, pp. 316-328.

[11] Creedy, J. and Gemmell, N. (2019) Illustrating income mobility: new measures. *Oxford Economic Papers*, 71, pp. 733-755.

[12] Creedy, J. and Gemmell, N. (2023) Summary measures of equalising mobility based on 'Three "I"s of Mobility' curves. *Journal of Income Distribution* (forthcoming).

[13] Creedy, J. and Sleeman, C. (2005) Adult equivalence scales, inequality and poverty. *New Zealand Economic Papers*, 39, pp. 51-83.

## About the Author

John Creedy is Professor of Public Finance at Wellington School of Business and Government, Victoria University of Wellington, New Zealand.
Email: john.creedy@vuw.ac.nz

Victoria University of Wellington
Te Herenga Waka
Chair in Public Finance

Working Papers in Public Finance