



A note on Inequality-Preserving Distributional Changes

John Creedy

WORKING PAPER 14/2014
October 2014

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:
<http://www.victoria.ac.nz/cpf/working-papers>

A Note on Inequality-Preserving Distributional Changes*

John Creedy

New Zealand Treasury and Victoria University of Wellington

Abstract

This note considers the problem of distributing a fixed amount of money ('income') among a given number of people, such that inequality (measured by either the Gini or Atkinson measure) takes a specified value. It is well known that simultaneous equations admit of many solutions where the number of variables exceeds that of equations (constraints). However, the approach examines cases where there are just one or two degrees of freedom, clarifying the resulting range of distributions. The properties of simultaneous disequalising and equalising transfers are discussed.

*In preparing this note I have benefited from encouragement by Peter Lambert and discussions with Chris Ball, Norman Gemmill and Justin van de Ven.

1 Introduction

Suppose it is required to distribute a fixed amount of money (for convenience, referred to as ‘income’) among a given number of people, such that the resulting Gini or Atkinson inequality measure takes a specified value.¹ Having specified this requirement, the question arises of whether it implies a unique income distribution. Is there more than one allocation that gives rise to the same inequality measure? From one perspective, this is a trivial problem. Imposing one or more moments of a distribution, along with an inequality measure, simply specifies linear and nonlinear constraints on the values in a distribution.² So long as the number of individuals exceeds the number of constraints, there are some degrees of freedom in selecting values: the simultaneous equations do not necessarily have a unique solution.

Despite this obvious property of simultaneous equations, the present note explores this question more formally, for Gini and Atkinson measures, in very simple contexts where the number of degrees of freedom is deliberately restricted to a minimum.³ Indeed, situations are explored where only one or two degrees of freedom are available. The results nevertheless are not restricted to the small populations, since it could simply be assumed that the other income levels are held fixed.⁴ Clearly, starting from some arbitrary distribution, an inequality measure can be preserved if an equalising transfer in one range of the distribution is appropriately matched by a disequalising transfer elsewhere. The problem is to determine the nature of such transfers if a continuous range of distributions is to be identified. For example, if two

¹Since these are relative measures, and are thus not affected by equal proportional changes in all incomes, the actual fixed amount is not important and the requirement could just as well be expressed in terms of income shares.

²As explained below, the Atkinson and Gini inequality measures are defined in each case as the proportional difference between an equally-distributed equivalent income, y_E , and the arithmetic mean. For the Atkinson measure, y_E is a power mean and for the Gini measure, it is a reverse-rank weighted mean.

³On these measures see, for example, Atkinson(1970) and Lambert (2001).

⁴The fact that incomes must be positive and, in the case of the Gini measure, the ranks of individuals play an important role, means that the choice of individual incomes is further restricted. These restrictions are not modelled explicitly here, but of course they are checked when investigating numerical examples and ultimately limit the range of alternative distributions.

measures, the mean and an inequality measure, must be constant, then the consideration of just three individuals (within a larger unspecified and fixed population) provides one degree of freedom and the ability to move between distributions using two income transfers operating in opposite directions.

Section 2 examines the Gini inequality measure. First, it considers the case where both the Gini and the arithmetic mean are held constant and it is found that a range of distributions can satisfy these constraints. Second, the variance is also held constant, and again a range of such distributions exists. Section 3 considers the case of the Atkinson inequality measure, showing that the allocation of a given income among a fixed number of people to achieve a specified Atkinson measure is, like the Gini, not unique. In each case illustrative examples are given of the types of transfer which can lead from one distribution to another having the same inequality.

2 The Gini Measure

This section examines the Gini measure, for which a number of different expressions can be found. One of the most frequently used formulae is the following. For a distribution, $x_1 < x_2 < \dots < x_n$, with arithmetic mean, \bar{x} , the standard Gini inequality measure can be written as:

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

The crucial term here is the sum, $\sum_{i=1}^n (n+1-i)x_i$, which, for constant G , n and \bar{x} , must be constant.⁵

2.1 Distributions with Equal Means

Another way of expressing the question raised in the introduction is to ask whether it is possible to have a mean-preserving change in the distribution of

⁵Writing the reverse-ranked weighted mean as $\bar{x}_R = \{\sum_{i=1}^n (n+1-i)x_i\} / \sum_{i=1}^n (n+1-i)$, it can be seen that, for large n , the Gini is effectively $1 - \bar{x}_R/\bar{x}$. Hence, \bar{x}_R is the equally distributed equivalent income for the Gini welfare function. However, the following discussion is for small n , and the form given above is not replication invariant.

x , for which the value of G remains unchanged? Would a judge, whose value judgements are characterised by those lying behind the Gini index, actually be indifferent between two (or possibly more) allocations?

The answer is particularly simple if $n = 2$ and there are only two individuals. If $x^* = x_2/2\bar{x}$ denotes the share of income going to the richer person, then (1) becomes⁶ $G = x^* - \frac{1}{2}$ and the allocation is:

$$x_2 = 2\bar{x}(G + 0.5) \quad (2)$$

with:

$$x_1 = 2\bar{x} - x_2 \quad (3)$$

Clearly, $x_2 > x_1$ and there is only one allocation satisfying the requirement.⁷ From this allocation, any change which leaves the arithmetic mean unchanged (for which $dx_1 = -dx_2$) causes G to change. Similarly, any change which leaves G unchanged requires a change in the arithmetic mean income. With two conditions (constant \bar{x} and G) and only two income levels, there are no degrees of freedom in selecting the allocation. This is of course not surprising with only two individuals and two constraints. Any possibility of having the two conditions satisfied by more than one allocation requires at least a degree of freedom in the choice of one of the x values.

Suppose $n = 3$. Then for unchanged mean and Gini, it is required to have, for given K and L :

$$3x_1 + 2x_2 + x_3 = K \quad (4)$$

and:

$$x_1 + x_2 + x_3 = L \quad (5)$$

From (5):

$$x_1 = L - x_2 - x_3 \quad (6)$$

and substituting in (4):

$$x_2 = 3L - K - 2x_3 \quad (7)$$

⁶This result for the above form of the Gini measure is given by Shorrocks (1995). An alternative version for the form of Gini that, for small n , lies between 0 and 1 is given by Subramanian (2002).

⁷In this case of $n = 2$, the maximum value that can be taken by G is 0.5.

Thus it is possible to set the value of x_3 , the higher income level, and allow the corresponding values of x_2 and x_1 to be easily obtained by using (7) and then (6). The above question thus reduces to finding, for given L and K whether more than one combination of x values can be obtained which preserves the rank order. Suppose $K = 10$ and $L = 6$, so that $\bar{x} = 2$. Setting $x_3 = 3$ produces $x_2 = 2$ and $x_1 = 1$. However, if $x_3 = 2.8$, the resulting values of x_2 and x_1 are 2.4 and 0.8. Both distributions give values of G of 0.222.⁸

It is well known that the Gini measure can be related to an area in the famous Lorenz curve diagram.⁹ The Gini measure is twice the area between the diagonal line of equality and the Lorenz curve. As an area measure, the Gini is not concerned with precisely which parts of the distribution contribute most to inequality. The two Lorenz curves are shown in Figure 1, where the curve for the second distribution [0.8, 2.4, 2.8] is the dashed line. This second distribution is closer to the line of equality for the higher part of the distribution but is further from the line of equality for the lower part of the distribution. Since the ranks of the individuals must remain unchanged, the two distributions involve simply moving the two Lorenz curve points (corresponding to the cumulative incomes associated with 1/3 and 2/3rds of the population) up or down somewhat. In each Lorenz curve the area contained by the Lorenz curve and the diagonal is the same. Intersecting Lorenz curves usually motivate the need for an explicit inequality measure, since a basic ‘dominance’ result does not apply, but in this case the Gini measure cannot distinguish between the two distributions.

Consider instead the distribution [0, 3, 3] for which the Lorenz curve follows the diagonal beyond the 2/3rds point and the Gini area is all contained to the left of that point (with $G = 0.333$). Another distribution with the same Gini is obtained simply by reducing inequality at the bottom of the dis-

⁸For these two distributions [1, 2, 3] and [0.8, 2.4, 2.8] the Atkinson values are quite different. For example, for ε of 0.2 and 1.2, the Atkinson values for the first distribution are 0.018 and 0.110 respectively, while for the second distribution they are 0.022 and 0.152.

⁹With incomes ranked in ascending order, the Lorenz curve plots, starting from the lowest income, the relationship between the cumulative proportion of total income and the proportion of people to whom it is attributed.

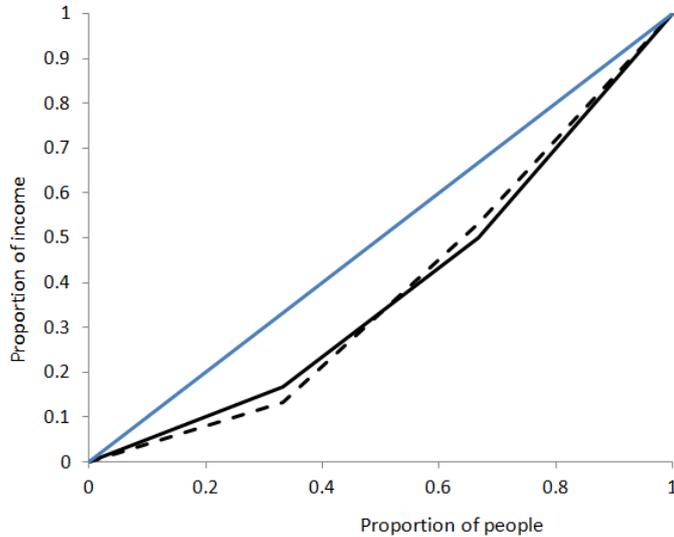


Figure 1: Two Lorenz Curves with Equal Ginis

tribution and increasing it at the top end: thus, to give just one alternative, $[0.2, 2.6, 3.2]$ is found to have the same Gini value and of course the same arithmetic mean. A way to view the two distributions is to see that, starting from $[0, 3, 3]$, an equalising transfer of 0.2 is made from person 2 to person 1, and at the same time a disequalising transfer of 0.2 is made from person 2 to person 3. There are two equal transfers from the middle person.

The nature of the effective transfers can be seen using the following approach, which begins from a given distribution and investigates the changes in x_1 , x_2 and x_3 for which the arithmetic mean and Gini are fixed. The Gini value for $n = 3$ can be written as:

$$G = 2 - \frac{2}{3} \left(\frac{3x_1 + 2x_2 + x_3}{x_1 + x_2 + x_3} \right) \quad (8)$$

Totalling differentiate G with respect to x_1 , x_2 and x_3 , and impose the condition $dx_1 = -(dx_2 + dx_3)$ to ensure that the arithmetic mean remains unchanged. Then setting $dG = 0$ gives the result, after a little algebra, that:

$$\frac{dx_2}{dx_3} = -2 \quad (9)$$

Hence the slope of the relevant constraint is fixed, irrespective of the x values.

Hence, this is precisely the change illustrated in the first example and the diagram above, where x_3 falls by 0.2 from 3 to 2.8 and x_2 rises by 0.4 from 2 to 2.4. And of course x_1 must change by $-(0.4 - 0.2) = -0.2$ and thus fall from 1 to 0.8. Importantly, there is thus a *range* of allocations satisfying the requirement that $G = 0.222$ and $\bar{x} = 2$. In each case, any number of new distributions can be found satisfying the required conditions. In summary, specifying an allocation of a fixed income among a given number of people, where the distribution must have a given Gini value, does not imply a unique allocation. Indeed, a *range* of distributions is consistent with the specified Gini.

2.2 A Common Mean and Variance

In the above example the variance of the first distribution is 0.667 while for the second distribution it is 0.747. The question arises of whether it is possible to have a change in the distribution which preserves both the mean and the variance, while also holding the Gini value constant. If a judge wishes to allocate a fixed sum of money to achieve both a given Gini and variance, is the resulting distribution unique? In this case, the additional constraint involves a loss of a further degree of freedom, so it is necessary to increase the population size to $n = 4$, the required conditions, for given values of K , L and M , are as follows. For constant Gini (for a given \bar{x}):

$$4x_1 + 3x_2 + 2x_3 + x_4 = K \quad (10)$$

and for a constant arithmetic mean:

$$x_1 + x_2 + x_3 + x_4 = L \quad (11)$$

Furthermore, a constant variance, for a given arithmetic mean, requires:

$$x_1^2 + x_2^2 + x_3^2 + x_4^2 = M \quad (12)$$

From (11):

$$x_1 = L - \sum_{i=2}^4 x_i \quad (13)$$

and substituting in (10):

$$\begin{aligned} x_2 &= 4(L - x_3 - x_4) + 2x_3 + x_4 - K \\ &= 4L - 2x_3 - 3x_4 - K \end{aligned} \quad (14)$$

Substituting (14) into (13) gives:

$$x_1 = K - 3L + x_3 + 2x_4 \quad (15)$$

Hence, substituting in (12) gives the quadratic:

$$\begin{aligned} M &= [L - x_3 - x_4 - \{4L - 2x_3 - 3x_4 - K\}]^2 \\ &\quad + \{4L - 2x_3 - 3x_4 - K\}^2 + x_3^2 + x_4^2 \end{aligned} \quad (16)$$

For given values of K , L and M , and setting a value of x_4 , the quadratic (16) can be solved for x_3 .¹⁰ The question then involves determining whether there are two real and distinct roots for x_3 and whether those resulting values give rise to corresponding values of x_2 and x_1 , using (13) and (14), which are positive and preserve the necessary rank order.

For example, suppose that $L = 8$, so that $\bar{x} = 2$, and $K = 15$, with $M = 21.54$. Setting $x_4 = 3.5$ it is found in this case that (16) has two roots of 2.2 and 2.8. Hence there are two distributions, $[0.2, 2.1, 2.2, 3.5]$ and $[0.8, 0.9, 2.8, 3.5]$ that have the same mean, variance and Gini value.¹¹ Going from one distribution to the other effectively involves two simultaneous transfers – one being disequalising while the other is equalising, from or towards person 2, depending on which is considered to be an ‘initial distribution’.

Consider the slope of the relevant constraint, relating x_3 and x_4 , by beginning with a particular distribution and investigating changes in x values for which the Gini and first two moments are constant. Totally differentiating:

$$G = \frac{5}{4} - \frac{1}{2} \left(\frac{4x_1 + 3x_2 + 2x_3 + x_4}{x_1 + x_2 + x_3 + x_4} \right) \quad (17)$$

¹⁰Extending the requirement to include a fixed third moment would thus be rather cumbersome as it would require the solution to a cubic in addition to a quadratic similar to (16).

¹¹In this case, $G = 0.3125$.

and using:

$$dx_1 = -(dx_2 + dx_3 + dx_4) \quad (18)$$

gives:

$$dx_2 = -2dx_3 - 3dx_4 \quad (19)$$

Then differentiating $x_1^2 + x_2^2 + x_3^2 + x_4^2$ totally, setting the result equal to zero and substituting for dx_1 and dx_2 using (18) and (19) gives:

$$\frac{dx_3}{dx_4} = \frac{x_3 - x_1 - 2(x_2 - x_1)}{x_4 - x_1 - 3(x_2 - x_1)} \quad (20)$$

Unlike the previous example, the slope of this constraint is not constant. Hence in general even small discrete changes from a given distribution do not satisfy the additional requirement.

Nevertheless the existence of a range of solutions can easily be seen by solving the above simultaneous equations for a different imposed value of x_4 . Starting from $[0.8, 0.9, 2.8, 3.5]$, suppose that there is an equalising transfer of 0.1 from person 4 (the richest individual) to person 3, while there is simultaneously a disequalising transfer of 0.1 from person 1 (the poorest) to person 2. This results in the distribution $[0.7, 1.0, 2.9, 3.4]$, which clearly ensures that total income is unchanged and is also found to maintain a fixed variance and Gini value. Hence, as before, only one degree of freedom is enough to generate the result that a range of distributions exist such that the first two moments and the Gini measure are constant. With a fifth person, giving two degrees of freedom, it is easy to see that a much wider range of distributions can satisfy the three constraints.¹²

3 The Atkinson Measure

This section considers the Atkinson inequality measure, A_ε , which is defined, for a relative inequality aversion parameter of ε , as:

$$A_\varepsilon = 1 - \frac{x_\varepsilon}{\bar{x}} \quad (21)$$

¹²For example, if a fifth person is added to $[0.2, 2.1, 2.2, 3.5]$, with $x_5 = 4$, then $L = 12$, $K = 27$, $\bar{x} = 2.4$ and $M = 37.5$. If one unit is transferred from person 5 to person 4, and at the same time one unit is transferred from person 2 to person 3, resulting in $[0.2, 2.0, 2.3, 3.6, 3.9]$, the three constraints are satisfied.

where x_ε is the equally-distributed equivalent income given by the power mean:

$$x_\varepsilon = \left\{ \frac{1}{n} \sum_{i=1}^n x_i^{1-\varepsilon} \right\}^{1/(1-\varepsilon)} \quad (22)$$

However, the question arises of whether two or more distributions can be found having the same Atkinson measure. Suppose only the mean is imposed, so that three values provide a necessary degree of freedom, and the sum, L , and equally distributed equivalent, x_ε , are given. Using $x_1 = L - x_2 - x_3$ and the definition of the equally distributed equivalent income, the value of x_2 , for given values of x_3 , L and x_ε , is given by the root or roots of:

$$(L - x_2 - x_3)^{1-\varepsilon} + x_2^{1-\varepsilon} + x_3^{1-\varepsilon} - 3x_\varepsilon^{1-\varepsilon} = 0 \quad (23)$$

It is found that this expression can have no real roots, one root or two real, positive and distinct roots. For example, suppose that $L = 6$ so that $\bar{x} = 2$, and for $\varepsilon = 0.3$, $x_\varepsilon = 1.8$. Setting $x_3 = 3$, it is found that there are two solutions for x_2 , equal to 0.2 and 2.8. However, the symmetry gives rise to corresponding values of x_1 of 2.8 and 0.2. Since the ranks are not relevant, having specified the value of x_3 , there is thus only one distribution, given by $[0.2, 2.8, 3.0]$, that is consistent with the imposed value of x_ε and the arithmetic mean, which together give A_ε .¹³ However, simply by setting x_3 to alternative values, a range of alternative distributions clearly exists for which the mean and Atkinson inequality measure are fixed. For example, starting from the above distribution, suppose that person 2 transfers 1.4 to person 3 and 0.4 to person 1. This combines an equalising with a disequalising transfer and results in the distribution $[0.6, 1.0, 4.4]$, which has the same Atkinson measure as the initial distribution.

With one extra individual, and hence an increase in the number of degrees of freedom to two, it is possible to generate an even wider range of possibilities. The equalising transfer in one range of the distribution can be

¹³Furthermore, using $dx_1 = -(dx_2 + dx_3)$ and setting the total differential, dA_ε , equal to zero gives $dx_2 = -dx_3(x_3 - x_1)/(x_2 - x_1)$. Since dx_2/dx_3 is not constant, this expression cannot be used to examine discrete changes. This contrasts with the Gini case discussed above where only the arithmetic mean and inequality are constant and there are three individuals.

combined with a disequalising transfer in another range of the distribution, involving a different pair of individuals (rather than the two transfers taking place from the middle-income person only). First, for $x_1 < x_2 < x_3 < x_4$, suppose a small amount dx^- is transferred from x_2 to x_1 while simultaneously dx^+ is transferred from x_3 to x_4 . Setting the total differential of x_ε equal to zero gives the result that, for an unchanged A_ε :

$$\frac{dx^+}{dx^-} = \frac{x_1^{-\varepsilon} - x_2^{-\varepsilon}}{x_3^{-\varepsilon} - x_4^{-\varepsilon}} \quad (24)$$

Again, this expresses the tangent to the required constraint imposed on the two changes to ensure that the Atkinson measure is constant. Given the nonlinearity involved, it cannot be used to consider a range of distributions.

Consider instead the simple extension of (23) to four individuals, giving:

$$(L - x_2 - x_3 - x_4)^{1-\varepsilon} + x_2^{1-\varepsilon} + x_3^{1-\varepsilon} + x_4^{1-\varepsilon} - 4x_\varepsilon^{1-\varepsilon} = 0 \quad (25)$$

It is now possible to set x_3 and x_4 and determine, for given L , ε and x_ε , the required values of x_1 and x_2 . As before, where two real roots for x_2 exist, these simply provide interchangeable values for x_1 and x_2 so that, since the ranks are not relevant, only one distribution satisfies the condition. However, two simultaneous transfers can be examined. Suppose that for $n = 4$ and $\varepsilon = 0.3$, it is required to have $\bar{x} = 2.0$ and $x_\varepsilon = 1.8$. First, setting $x_4 = 4$ and $x_3 = 2.5$ gives the result that $x_2 = 1.35$ and $x_1 = 0.15$. Then suppose that there is a disequalising transfer at the top end of the distribution such that x_4 becomes 4.2 and x_3 becomes 2.3. Solving (25) gives new values for $x_2 = 1.3$ and $x_1 = 0.2$, for which the same conditions are satisfied for L and x_ε . Hence an equalising transfer at the bottom end of the distribution of 0.05 from x_2 to x_1 can be combined with a disequalising transfer at the top end, of 0.2 from x_3 to x_4 , to maintain the Atkinson inequality measure.¹⁴

¹⁴For this case of $dx^- = 0.05$, the expression in (24) gives a value of $dx^+ = 0.43$, showing the extent of the nonlinearity of the constraint, since for the discrete change examined, it is necessary only to transfer 0.2 from x_3 to x_4 .

4 Conclusions

This note has considered the problem of distributing a fixed amount of money ('income') among a given number of people, such that inequality, as measured by either the Gini or Atkinson measure, takes a specified value. In particular, the question arises of whether a unique distribution is implied. In general, it is trivially obvious that, given a set of constraints (in the form of moments of the distribution), a range of solutions to the implied simultaneous equations can generally be found. However, this note has explored the characteristics of the two inequality measures for cases where only one or two degrees of freedom are available. The examination of small populations is not a restriction, since the income values other than the small number being considered may simply be regarded as being fixed. The measures differ in that the inequality constraint for the Gini involves a linear combination of income values, involving the rank positions, whereas for the Atkinson measure a nonlinear constraint is imposed, involving the degree of relative inequality aversion.

For both inequality measures, the intuition was confirmed that only one degree of freedom is necessary to generate a range of possible distributions, for cases where the mean and inequality are constrained and (in the case of the Gini measure) where the first two moments and inequality are constrained. For the Gini measure it was found that the choice of one income level (the degree of freedom) can give rise to two different distributions giving the same Gini value (and moments): that is, the simultaneous equations can generate two distinct solutions (satisfying the requirement that values are non-negative and preserve the rank order). Furthermore, a range of distributions is obtained simply by varying the income level imposed. In the case of the Atkinson measure, when one of the income levels is imposed (using the degree of freedom), the simultaneous equations can give rise to only one feasible solution: where two solutions are generated, they imply the same distribution since the rank order of individuals is not important. Nevertheless, a range of distributions is consistent with a fixed Atkinson measure, obtained by varying the imposed income level. In both Gini and Atkinson cases, the variations are clearly identified with simultaneous equalising and disequal-

ising transfers. For a sufficiently large population, it should in general be possible to find multiple allocations for which any number of moments of the distribution, along with an inequality measure, are equal.¹⁵

The important question arises of whether this feature matters for the practical measurement of inequality. To the extent that the two measures can be interpreted as reflecting explicit (though quite different) sets of value judgements of a disinterested judge, the fact that a ‘stable’ value of inequality – along with an unchanged arithmetic mean – is consistent with quite substantial changes in the precise nature of the frequency distribution of income is not important. The judge is by definition indifferent to all such distributions. Hence, the main implications are for the reporting of inequality measures, where exercises involve examination of the implications of adopting alternative value judgements.¹⁶ Thus readers of a document showing that a particular summary inequality index is constant, and who do not necessarily share the value judgements implicit in the measure used, need to be well aware that the distribution may have changed in ways that those readers consider to be important. This again supports the presentation of a range of summary measures, or at least the use of extensive sensitivity analyses, even though – when broad agreement is obtained – space constraints may mean that all results cannot be reported.

¹⁵To give just one example using numerical methods, the distributions [1, 2, 4, 5, 10] and [0.8, 2.6, 6.3, 5.2, 10] have the same Gini and first three moments. The second can be obtained from the first by using two disequalising transfers and one equalising transfers. Small variations in the top income (of 10) generate two solutions for each imposed value of x_5 .

¹⁶Of course, in many reports inequality measures are reported with no clarification of the value judgements involved. The Gini, for example, is often described simply as being ‘widely used’, as if that were sufficient justification for the reader to accept the results without question.

References

- [1] Atkinson, A.B. (1970) On the measurement of inequality. *Journal of Public Economics*, 2, pp. 244-263
- [2] Lambert, P.J. (2001) *The Distribution and Redistribution of Income*. Manchester: Manchester University Press.
- [3] Shorrocks, A.F. (2005) Inequality values and unequal shares. UNU-WIDER, Helsinki.
- [4] Subramanian, S. (2002) An elementary interpretation of the Gini inequality index. *Theory and Decision*, 52, pp. 375-379.

About the Authors

John Creedy is Professor of Economics at Victoria Business School, Victoria University of Wellington, New Zealand, and a Principal Advisor at the New Zealand Treasury.

Email: john.creedy@vuw.ac.nz

