

Interpreting Inequality Measures and Changes in Inequality

John Creedy

New Zealand Treasury Working Paper 14/20

November 2014

DISCLAIMER

The views, opinions, findings, and conclusions or recommendations expressed in this Working Paper are strictly those of the author(s). They do not necessarily reflect the views of the New Zealand Treasury or the New Zealand Government. The New Zealand Treasury and the New Zealand Government take no responsibility for any errors or omissions in, or for the correctness of, the information contained in these working papers. The paper is presented not as policy, but with a view to inform and stimulate wider debate.



THE TREASURY
Kaitohutohu Kaupapa Rawa

New Zealand Government

NZ TREASURY
WORKING PAPER
14/20

Interpreting Inequality Measures and Changes in Inequality

MONTH/YEAR

November 2014

AUTHORS

John Creedy
Principal Advisor/Professor of Public Economics
New Zealand Treasury/Victoria University of Wellington
No. 1 The Terrace
Wellington
New Zealand
Email: John.Creedy@treasury.govt.nz
Telephone: ++64 4 917 6893

ISBN (ONLINE)

978-0-478-43604-4

URL

Treasury website at November 2014:
<http://www.treasury.govt.nz/publications/research-policy/wp/2014/14-20/twp14-20.pdf>
Persistent URL: <http://purl.oclc.org/nzt/p-1701>

ACKNOWLEDGEMENTS

I have benefited from detailed comments by Anita King and Justin van de Ven on an earlier draft. I am grateful to Nicolas Hérault and Stephen Jenkins for drawing my attention to literature on the pivotal income.

NZ TREASURY

New Zealand Treasury
PO Box 3724
Wellington 6008
NEW ZEALAND
Email: information@treasury.govt.nz
Telephone: +64 4 472 2733
Website: www.treasury.govt.nz

Abstract

This paper explores, in the context of the Atkinson inequality measure, attempts to make interpretations of orders of magnitude transparent. One suggestion is that the analogy of sharing a cake among a very small number of people provides a useful intuitive description for people who want some idea of what an inequality measure ‘actually means’. In contrast with the Gini measure, for which a simple ‘cake-sharing’ result is available, the Atkinson measure requires a nonlinear equation to be solved. Comparisons of ‘excess shares’ (the share obtained by the richer person in excess of the arithmetic mean) for a range of assumptions are provided. The implications for the ‘leaky bucket’ experiments are also examined. An additional approach is to obtain the ‘pivotal income’, above which a small increase for any individual increases inequality. The properties of this measure for the Atkinson index are also explored.

JEL Classification: D331; D63

Keywords: Inequality; Atkinson Measure; Excess Share.

Executive Summary

This paper considers the problem of conveying just what a particular income inequality value, or a change in that value, might mean in ways that are easy to understand. Emphasis is placed on the widely-used Atkinson inequality measure, which is based on explicit value judgements involving a parameter reflecting the degree of inequality aversion of an independent judge.

In view of the strong interest in inequality and changes in inequality over time, and the widespread reference to inequality measures, both in evaluating policy changes and in making a case for particular types of policy change, there is a strong incentive to try to provide empirical results that provide transparent guidance about orders of magnitude.

One approach is based directly on the trade-off between equity and efficiency that is implied by the welfare, or evaluation, function (summarising a judge's value judgements) on which the inequality measure is based. For example, given an annual growth rate, it could be said that a judge is prepared to give up a certain number of years of growth to obtain a given percentage reduction in inequality.

The second approach takes as its starting point the suggestion that considering the division of a cake between a small number of people provides a clear idea of the extent of relative inequality. Hence, an inequality measure can be reduced to an equivalent expression of just two income (share) values. Earlier studies have explored this approach for the Gini inequality measure, for which a very simple result is available. It is shown here that, for the Atkinson measure, some interesting insights are available, but it is not clear that the approach can provide the kind of information that non-specialists can easily digest in comparing different values of inequality.

The third approach involves the concept of a pivotal income value. This is the income level such that an increase in any income above the pivotal value causes the inequality measure to increase. This is seen to be very easily calculated given the Atkinson measure and its associated degree of inequality aversion.

It is suggested that the information provided by converting inequality measures to comparisons involving a simplified cake-cutting exercise, or concentrating on a pivotal income, cannot be expected to provide easy and clear answers to the communication challenge. But with careful use they may provide useful supplementary descriptions. Such approaches do at least preserve the basic

value judgements behind the inequality measures and allow for the implications of alternative value judgements to be investigated. This is preferred to the kind of oversimplified comparisons that disguise the implicit value judgements of those reporting results.

Contents

Abstract	i
Executive Summary	ii
1 Introduction	1
2 Previous Results for the Gini Measure	3
3 The Atkinson Inequality Measure	4
4 The Abbreviated Welfare Function	5
5 An Equivalent Small Distribution	6
5.1 n-Person Shares	6
5.2 The Leaky Bucket Experiment	9
6 The Pivotal Income	11
7 Conclusions	14

List of Figures

Figure 1 – Solving for the Two-Income Case: Atkinson Inequality Index of 0.25 and Inequality Aversion of 0.2	7
Figure 2 – Variation in x with A: Epsilon = 0.2	8
Figure 3 – Variation in x with A: Epsilon = 0.9	9
Figure 4 – Ratio of Pivotal Income to Arithmetic Mean and Variation in Inequality Aversion	12
Figure 5 – Variation in Pivotal Income with Inequality	13
Figure 6 – Proportion Below Pivotal Income and Inequality: Lognormal Distribution	13

List of Tables

Table 1 – Top Income, x, for Alternative Inequality Measures and Population Size . .	7
Table 2 – Excess Shares for Alternative Inequality Measures and Population Size .	8
Table 3 – Tolerance for Leaky Bucket as Percentage of Amount Taken: n = 6	10
Table 4 – Atkinson Inequality Measure for given x and Alternative Aversion Coefficients: n = 6	10

Interpreting Inequality Measures and Changes in Inequality

1 Introduction

Anyone who reports a summary inequality measure and its changes over time, or who provides a distributional evaluation of a potential policy change in terms of a change in inequality, is sure to be asked the difficult question, ‘what does this mean?’ Is, say, a five per cent increase in an inequality measure a major concern for an inequality-averse judge, or may it be regarded as quite small? Whereas it is relatively easy to form a view about a change in an aggregate income measure – a simple sum of values – it is intuitively much more difficult to appreciate orders of magnitude of an inequality measure which is (usually) based on nonlinear transformations of many individual values. Of course, the Lorenz curve provides an extremely convenient graphical device for illustrating aspects of relative inequality.¹ Yet in many practical contexts it is necessary to provide a quantitative indication of inequality and its changes. The aim of this paper is to explore several approaches that have been proposed to help interpret inequality and its changes in the context of the Gini measure. These approaches are applied to the Atkinson inequality measure.

One possible approach, proposed by Subramanian (2002) and extended by Shorrocks (2005), is to consider an alternative artificial distribution that has the same value of the inequality measure as the actual distribution but consists of only two income levels (or, equivalently, income shares), though there may be more than two individuals in the artificial distribution. A basic analogy with a ‘cake-cutting exercise’ is therefore involved. Both Subramanian and Shorrocks concentrated mainly on the Gini inequality measure. It is necessary to consider in further detail whether Subramanian’s claim, followed by Shorrocks, that ‘the clearest idea we can have of the extent of relative inequality’ is unambiguously provided by the share received by the poorer person (or persons) in a finite-sample comparison. Does the analogy of sharing a cake among a very small number of people actually provide the kind of intuitive description that is useful for people who want some idea of what an Atkinson inequality measure and its changes ‘actually mean’?

¹ Furthermore, its relationship with value judgements, such as the principle of transfers, has been well explored. The concepts of Lorenz Dominance and Generalised Lorenz Dominance play a central role in inequality and welfare comparisons; see, for example, Lambert (2001).

The simple observation that a small income increase for a high-income person leads to an increase in inequality, but an increase for a low-income person results in a reduction in inequality, gives rise to the idea of a particular income level that in some sense ‘divides’ the low and high incomes. This idea appears to have given rise to three independent analyses. First, Hoffmann (2001) referred to the dividing line as the ‘relative poverty line’ and explored its value for a range of inequality measures (including the Gini and the generalised entropy class of measures).² Second, Lambert and Lanza (2006) defined the dividing line as a ‘benchmark income or position’, obtaining a wide range of results.³ Third, Corvalan (2014), used the concept of a ‘pivotal’ individual, defined as the person such that if an additional small income increase is given to a poorer individual, the inequality index falls. Corvalan concentrated on the Gini measure of inequality. The additional insights provided by this approach are explored here in the context of the Atkinson index.

First, Section 2 briefly summarises the main results regarding the Gini measure. Section 3 provides a brief reminder of the definition of the Atkinson measure. The judge’s evaluation function that lies behind the Atkinson measure allows for a well-defined trade-off between inequality and total income (or, as it is sometimes expressed, between ‘equity and efficiency’). This trade-off therefore provides an initial indication of what inequality ‘means’: it makes it possible to talk about the extent of total income growth that would be foregone by a judge in order to achieve a given reduction in inequality. This is described in Section 4.

The Subramanian–Shorrocks approach applied to the Atkinson measure is then discussed in Section 5. Examples are given of the variation in the ‘excess share’, for different finite-sample sizes, as inequality changes. The fact that only two income values (even where $n > 2$) are involved means that it is also possible to exploit the ‘leaky bucket’ thought experiment. This is the device Atkinson introduced in order to help clarify the specification of orders of magnitude regarding an ‘inequality aversion’ parameter. The implications of applying this thought experiment to the n -person shares are also investigated in Section 5. Section 6 then turns to the ‘pivotal income’ associated with any degree of inequality aversion and value of the Atkinson inequality measure. Brief conclusions are in Section 7.

² He also reported values of the relative poverty line for five regions of Brazil.

³ They also considered the role of the benchmark income in the leaky bucket experiment, in the context of effects on inequality measures, rather than the usual reference to constant ‘social welfare’ evaluations.

2 Previous Results for the Gini Measure

The first approach, along ‘cake-cutting’ lines and in the context of the Gini inequality measure, was made by Subramanian (2002). He considered a distribution having the same Gini value as the actual distribution but consisting of only two income levels (or, equivalently, income shares). The restriction to just two incomes enables some simple comparisons to be made. He suggested that, ‘our intuitive comprehension of the notion of inequality is sharpest in the context of the canonical two-person cake-sharing problem’ (2002, p. 375). He argued that, ‘the share σ of the cake going to the poorer person furnishes just about the clearest idea we can have of the extent of relative inequality’. Subramanian used a definition of the Gini measure, G_s , that is normalised to the range $0 \leq G_s \leq 1$ for finite sample sizes, and showed that $G_s = 1 - 2\sigma$.

Shorrocks (2005) subsequently showed that the standard Gini measure, G , provides a particularly simple illustration with a more immediate interpretation.⁴ In the two-person case equality is obtained when each person has 50 per cent of the total income. If the richer person has a share of x , the poorer person has a share of $1 - x$: these shares may also be considered as income levels if total income is normalised to one. Shorrocks showed that G is equal simply to $x - 0.5$: he referred to this as the ‘excess share’, since it is the difference between x and the equal share (which he called the ‘fair share’). A complication is that the two-person standard Gini has a maximum value of 0.5, so this case can only be considered where the actual Gini under discussion is less than 0.5 and, even then, the comparison produces a rather wide difference between the two income levels, even for modest Gini values. The approach was extended to allow for more than two individuals in receipt of the lower income value. Shorrocks demonstrated that for a finite population of n individuals, with the richer person having x and each of the $n - 1$ poorer people receiving $(1 - x) / (n - 1)$, the Gini measure is $G = x - 1/n$ and, since equality implies shares of $1/n$, the Gini continues to be interpreted as the ‘excess share’.⁵

Shorrocks mentioned that such shares can be produced for any inequality measure, though none is as simple as the Gini. He compared decile shares and excess shares for a number of countries and inequality measures, including the Atkinson measure (though he did not explore its properties, giving only the formula for the

⁴ The standard Gini measure is ‘replication invariant’, unlike G_s .

⁵ This is demonstrated in Appendix A below.

index), and pointed out that in some cases the ranking by inequality measure is not the same as by excess share (for different n values).

As mentioned above, the approach adopted by Corvalan (2014) retained the actual distribution but involved the concept of a ‘pivotal’ individual. For the Gini measure, which of course is based on the ranking of individuals’ incomes, Corvalan obtained a simple expression for the rank of the pivotal person. It is thus possible to obtain, given the Gini measure, the percentile above which a (rank preserving) increase in income for any individual produces an increase in inequality (and thus below which an income increase produces a reduction in inequality). For a large population, Corvalan (2014, p. 600) showed that the percentile above which a transfer increases inequality is given simply by $100(1 + G)/2$. Corvalan did not explore the implications for the Atkinson inequality measure.

3 The Atkinson Inequality Measure

Consider a population of n individuals with incomes of y_1, y_2, \dots, y_n . Suppose an independent judge evaluates income distributions using the following Social Welfare Function:

$$W = \left(\frac{1}{n}\right) \sum_{i=1}^n \frac{y_i^{1-\varepsilon}}{1-\varepsilon} \quad (1)$$

where ε reflects the degree of relative inequality aversion of the judge. This welfare function is individualistic, additive, Paretean and satisfies the ‘Principle of Transfers’ (such that a transfer from a richer to a poorer person, which does not affect their relative positions, represents an ‘improvement’). Define y_ε as the ‘equally distributed equivalent’ income level. That is, y_ε is the income which, if received by all n individuals, gives the same value of W as the actual distribution. Hence y_ε is the power mean given by:

$$y_\varepsilon = \left\{ \frac{1}{n} \sum_{i=1}^n y_i^{1-\varepsilon} \right\}^{1/(1-\varepsilon)} \quad (2)$$

Atkinson’s (1970) inequality measure, A , is defined as:

$$A = 1 - \frac{y_\varepsilon}{\bar{y}} \quad (3)$$

where \bar{y} is arithmetic mean income. Hence inequality is expressed in terms of a ratio of two measures of location (or central tendency), here the ratio of a power mean of order, $1 - \varepsilon$, to the arithmetic mean, \bar{y} . An advantage of the Atkinson measure is thus its clear connection to the value judgements involved in the choice

of (1).⁶ Nevertheless, interpreting both the size of ε and a change in A are far from transparent.

4 The Abbreviated Welfare Function

One way to interpret a change in inequality is to refer to the associated abbreviated welfare function which reflects the trade-off between total income and its equality. This trade-off is implicit in the value judgements behind the welfare, or evaluation, function itself. For the Atkinson inequality measure, the abbreviated function, denoted W^* , is:

$$W^* = \bar{y}(1 - A) \quad (4)$$

Thus W^* shows (loosely speaking) the trade-off between ‘equity and efficiency’ that is implicit in W and the inequality measure, A , defined above, when equity, E , is measured by $E = 1 - A$. Hence any change in inequality can be expressed, in social welfare function terms, as equivalent to a particular change in \bar{y} . Totally differentiating (4) gives:

$$\frac{dW^*}{W^*} = \frac{d\bar{y}}{\bar{y}} + \frac{dE}{E} = \frac{d\bar{y}}{\bar{y}} + \frac{d(1 - A)}{(1 - A)} \quad (5)$$

and the trade-off (for which $\frac{dW^*}{W^*} = 0$) is such that a 1 percentage change in equity is equivalent to a 1 percentage change in \bar{y} , and:

$$\left. \frac{d\bar{y}}{\bar{y}} \right|_{W^*} = -\frac{d(1 - A)}{(1 - A)} = \left(\frac{A}{1 - A} \right) \frac{dA}{A} \quad (6)$$

This expression shows that the percentage change in \bar{y} that is equivalent to a given percentage change in inequality, A , is linearly (proportionally) related to the latter, with a slope that depends on $A/(1 - A)$. Obviously if $A = 0.5$, there is a one-to-one relationship between the percentage changes. With $A < 0.5$, for increasing inequality $\left. \frac{d\bar{y}}{\bar{y}} \right|_{W^*} < \frac{dA}{A}$, and vice versa for $A > 0.5$. For example, an increase in A from 0.25 to 0.30 represents a 20 per cent increase, which is equivalent in social welfare function terms to a reduction in arithmetic (or total) income of 6.67 per cent. An increase in A from a higher value of 0.30 to 0.36, which is also a 20 per cent increase in the Atkinson measure is, however, equivalent in social welfare function terms to a reduction in \bar{y} of 8.57 per cent.

⁶ As shown in the Appendix, Subramanian (2002) related the Gini measure to a welfare function based on rank-order weights, using a corresponding equally distributed equivalent income level.

5 An Equivalent Small Distribution

It is difficult to envisage what a particular value of A , and a change in A , implies in view of the fact that incomes are usually distributed among a very large number of people. One approach to interpreting orders of magnitude, mentioned in the introduction, is effectively to consider an artificial small population having the same inequality measure. The approach asks what a distribution with just two income levels would look like, having the same A value and of course evaluated for the same value of ε . However, such a division must necessarily polarise the distribution and, as seen below, can therefore provide an exaggerated picture of what an inequality measure ‘means’.

5.1 n-Person Shares

Suppose there are just two types of individual and, for convenience only, the total income is normalised to one unit. As relative inequality is of concern, the income values can also be considered as income shares. One person has an income of x , and the remaining $n - 1$ people have an income of $(1 - x) / (n - 1)$ each, and arithmetic mean income is $1/n$. Thus:


$$A = 1 - n \left[\frac{1}{n} \left\{ x^{1-\varepsilon} + (n-1) \left(\frac{1-x}{n-1} \right)^{1-\varepsilon} \right\} \right]^{1/(1-\varepsilon)} \quad (7)$$

Given a value of A obtained from an actual income distribution, the aim is to solve (7) for x . Rearranging this, it can be seen that x is given by the root (or roots) of:

$$x^{1-\varepsilon} + (n-1)^\varepsilon (1-x)^{1-\varepsilon} - n \left(\frac{1-A}{n} \right)^{1-\varepsilon} = 0 \quad (8)$$

It may initially seem natural to consider just two individuals in the artificial population, producing only a richer and a poorer person and thereby making the properties more transparent. However, as with the standard Gini measure, it is not always possible to solve (8), remembering also that feasible solutions require $0 < x < 1$. Allowing for $n > 2$ can enable a feasible solution to be obtained. This aspect can be seen by considering the case where $A = 0.25$ for an inequality aversion parameter of $\varepsilon = 0.2$. Figure 1 shows the variation in the left hand side of (8) as x is varied over the feasible range, for three different values of n . Clearly with $n = 2$ and $n = 4$ there are insufficient poorer people to generate enough inequality with just the two income levels. When $n = 6$, a feasible solution can be obtained; the profile intersects the horizontal axis for a high value of x .

Figure 1: Solving for the Two-Income Case: Atkinson Inequality Index of 0.25 and Inequality Aversion of 0.2



Consider how the values of x vary as the Atkinson measure increases. Table 1 illustrates the effect of increasing A by 20 per cent, from an initial value of $A = 0.25$. A dash (–) in the table indicates that for the combination of n and ε , no feasible solution for x exists: it is not possible to generate the required inequality level with such a small number of individuals. The table shows that for the large (20 per cent) increase in A , the associated increase in the income (share) of the rich person in the small-sample construction varies depending on the sample size, n , and the degree of inequality aversion. For $\varepsilon = 0.5$ and $n = 2$, the value of x increases by 3.2 per cent as A increases by 20 per cent. This increases to 6.7 per cent the same ε but for $n = 4$. For the low inequality aversion parameter of $\varepsilon = 0.2$ in combination with $n = 6$ and $n = 8$, x must increase by 6.7 and 7.0 per cent respectively. Yet for $\varepsilon = 0.9$ in combination with $n = 6$ and $n = 8$, x must increase by 7.2 and 17.4 per cent respectively.

Table 1: Top Income, x , for Alternative Inequality Measures and Population Size

ε	$A = 0.25$				$A = 0.30$			
	n				n			
	2	4	6	8	2	4	6	8
0.2	–	–	0.90	0.85	–	–	0.96	0.91
0.5	0.93	0.75	0.66	0.60	0.96	0.80	0.71	0.65
0.9	0.84	0.63	0.55	0.46	0.87	0.68	0.59	0.54

Table 2 converts the values of x in Table 1 into their corresponding ‘excess shares’, $x - 1/n$. In the Gini case these would, as Shorrocks (2005) showed, all equal the appropriate standard Gini measure, but this simple property does not carry over to the Atkinson measure. One feature of these values is that, for given ε , the variation in the excess share as n increases is not monotonic. It increases initially as n increases from $n = 2$, and then declines slightly after $n = 4$. However,

the excess share decreases monotonically as the degree of inequality aversion, ε , increases, for which more weight is attached to the lower income range. For a given value of n , the richer share, x , must decrease as ε increases, because it is easier to achieve the given inequality measure, A , with a lower share when aversion to inequality is higher. The excess share is by definition equal to $x - \frac{1}{n}$, so this must also fall as ε increases, given that such comparisons hold n constant.⁷ The changes in the excess share vary as A increases from 0.25 to 0.30: it varies from about 8 per cent to over 20 per cent (for the $\varepsilon = 0.9$ combined with $n = 8$ case).

Table 2: Excess Shares for Alternative Inequality Measures and Population Size


ε	$A = 0.25$				$A = 0.30$			
	n				n			
	2	4	6	8	2	4	6	8
0.2	–	–	0.73	0.73	–	–	0.79	0.79
0.5	0.43	0.50	0.49	0.48	0.46	0.55	0.54	0.53
0.9	0.34	0.38	0.38	0.34	0.37	0.43	0.42	0.42

The variation in x as A varies, for given n and ε , can be examined more formally as follows. Implicit differentiation of (8) shows that the slope of this relationship is given by:

$$\frac{dx}{dA} = \frac{n \left(\frac{1-A}{n}\right)^{-\varepsilon}}{(n-1)^\varepsilon (1-x)^{-\varepsilon} - x^{-\varepsilon}} \quad (9)$$


Although this relationship is clearly nonlinear, it can be seen from Figures 2 and 3 that they are approximately linear. The two figures show the variation for values of ε of 0.2 and 0.9 respectively. In the case of Figures 2 no profile is shown for $n = 2$ because there is no feasible root of (8) for all values of A shown.

Figure 2: Variation in x with A : Epsilon = 0.2



⁷ Curiously, Shorrocks reports (2005, Table 3) values of the excess share which, for $n = 10$, increase as ε is increased.

Figure 3: Variation in x with A : Epsilon = 0.9



5.2 The Leaky Bucket Experiment

The restriction to two income levels in the finite population construct means that the 'leaky bucket' thought experiment can be applied directly to those two values. In making comparisons between two different values of inequality it is useful to keep in mind what is implied in terms of the willingness to tolerate leaks in making an income transfer between the two incomes, that is from the richer person to each of the other $n - 1$ people who share a common income level.

From (1) consider changes in W resulting from a change in the two income levels y_1 and y_2 , with $y_1 > y_2$. Totally differentiating gives:

$$ndW = y_1^{-\varepsilon} dy_1 + (n - 1) y_2^{-\varepsilon} dy_2 \quad (10)$$

Hence, for constant W :

$$\left. \frac{dy_2}{dy_1} \right|_W = - \left(\frac{1}{n - 1} \right) \left(\frac{y_2}{y_1} \right)^{\varepsilon} \quad (11)$$

If Δy_1 is taken from the richer person, the amount that needs to be given to each of the $n - 1$ people with y_2 to hold W constant, Δy_2 , is given by:

$$\Delta y_2 = \Delta y_1 \left(\frac{1}{n - 1} \right) \left(\frac{y_2}{y_1} \right)^{\varepsilon} \quad (12)$$

Thus the minimum proportionate increase in y_2 is given by:

$$\frac{\Delta y_2}{y_2} = \frac{\Delta y_1}{y_1} \left(\frac{1}{n - 1} \right) \left(\frac{y_1}{y_2} \right)^{1-\varepsilon} \quad (13)$$

The maximum leak tolerated by the independent judge, expressed as a proportion of the amount taken from y_1 , is thus:

$$\frac{\Delta y_1 - (n - 1) \Delta y_2}{\Delta y_1} = 1 - \left(\frac{y_2}{y_1} \right)^{\varepsilon} \quad (14)$$

Table 3 illustrates the implications of different inequality aversion parameters for the tolerance of leaks in taking a small amount from the richer person with x , and giving equal amounts to each of the poorer individuals such that social welfare is unchanged. All the cases illustrated in the table are for $n = 6$, and comparisons are given for two inequality values. Hence the top left-hand cell in the main body of the table shows that if $A = 0.25$ when $\varepsilon = 0.2$, the judge tolerates a leak of 53.5 per cent of the amount taken from the richer person, with x given by the corresponding value in Table 1. Moving along the row, if the associated value of ε is 0.9, again giving $A = 0.25$, the judge would tolerate a leak of 80.2 per cent of the amount taken from the richer person. Corresponding amounts for the higher inequality value are shown in the right-hand block of the table which, as expected, reveal a greater tolerance for leaks.

Table 3: Tolerance for Leaky Bucket as Percentage of Amount Taken: $n = 6$

$A = 0.25$			$A = 0.30$		
$\varepsilon = 0.2$	$\varepsilon = 0.5$	$\varepsilon = 0.9$	$\varepsilon = 0.2$	$\varepsilon = 0.5$	$\varepsilon = 0.9$
53.5	67.6	80.2	61.2	71.3	83.2

Table 4 illustrates a different kind of sensitivity. This shows the values of A that are obtained for the x value corresponding to a specified A and ε , but for different ε values. The bold diagonal values are the assumed Atkinson measures of 0.25 and 0.30. Consider the first row, for $\varepsilon = 0.2$. The value of x that is associated with this combination ($n = 6$, $A = 0.25$ and $\varepsilon = 0.2$) is, from Table 1, equal to 0.9. If the resulting values of 0.90 for the richer person and $0.10/5$ for each of the poorer people are then used to calculate inequality, based on a different inequality aversion parameter, of $\varepsilon = 0.5$, the Atkinson measure would be 0.55. Similarly if $\varepsilon = 0.9$ were applied to those incomes, an inequality measure of 0.75 would result. Looking along the last three rows of the table therefore shows that the values of inequality are indeed highly sensitive to variations in ε , starting from the income shares that are consistent with the relevant value (from the row) of ε .

Table 4: Atkinson Inequality Measure for given x and Alternative Aversion Coefficients: $n = 6$

	$\varepsilon = 0.2$	$\varepsilon = 0.5$	$\varepsilon = 0.9$	$\varepsilon = 0.2$	$\varepsilon = 0.5$	$\varepsilon = 0.9$
$\varepsilon = 0.2$	0.25	0.55	0.75	0.30	0.66	0.87
$\varepsilon = 0.5$	0.11	0.25	0.38	0.14	0.30	0.44
$\varepsilon = 0.9$	0.07	0.16	0.25	0.09	0.19	0.30

Atkinson's (1970) initial discussion of the leaky bucket experiment was in the context of finding a way to obtain a clearer view of the meaning of different values of ε , using transfers between just two individuals with incomes of y_1 and y_2 , say. This produces the result that, after taking $\Delta y_1 = 1$ from person 1 (the richer

person), it is necessary to transfer only Δy_2 to person 2 (the poorer person) where:

$$\Delta y_2 = \left(\frac{y_2}{y_1} \right)^\varepsilon \quad (15)$$

Comparisons can then be made for various values of the ratio of incomes. Thus if $\frac{y_2}{y_1} = \frac{1}{2}$, aversion of $\varepsilon = 0.2$ suggests that a leak of 0.13 is tolerated (13 per cent of the dollar taken from person 1). This suggests that 0.2 is a relatively low degree of aversion to inequality. Yet, when a distribution is ‘compressed’ into two income levels (shares of 1 unit), the resulting spread of incomes produces a tolerance, as shown in Table 3, that is very much greater, for the same aversion coefficient. This is because, as mentioned earlier, the need to generate a given inequality value with just two income values (shares) means that those values must be ‘pushed’ very far apart. The ratio of the two incomes is very much greater than the usual range considered when the interpretation of ε is being discussed. For the same case of $\frac{y_2}{y_1} = \frac{1}{2}$, a high aversion of 1.5 suggests a tolerance for a leak of 35 per cent of the dollar taken from the richer person. This is considerably lower than any of the values shown in Table 3.

These examples show that the transformation or abbreviation of a distribution to only two income values (with a small number of individuals) means that those incomes have to be polarised to a degree that is unrealistic and well beyond the range that people generally consider when thinking about inequality aversion.

6 The Pivotal Income

Consider the concept of the pivotal income, defined in the introduction as the income below which a small increase leads to a reduction in inequality, applied to the Atkinson inequality measure. This measure is differentiable, so that the pivotal income, y^* , is given by the solution to $\partial A / \partial y_i = 0$.⁸ Thus:

$$\frac{\partial A}{\partial y_i} = -\frac{y_\varepsilon y_i^{-\varepsilon}}{n\bar{y}} \frac{1}{\frac{1}{n} \left(\sum_{i=1}^n y_i^{1-\varepsilon} \right)} + \frac{y_\varepsilon}{n\bar{y}^2} \quad (16)$$

Setting this equal to zero and rearranging gives:

$$(y^*)^{-\varepsilon} = \frac{\frac{1}{n} \sum_{i=1}^n y_i^{1-\varepsilon}}{\bar{y}} \quad (17)$$

From (2) and (3):

$$\frac{1}{n} \sum_{i=1}^n y_i^{1-\varepsilon} = ((1 - A)\bar{y})^{1-\varepsilon} \quad (18)$$

⁸ The following result is simply a special case of the much more general results obtained for ‘non-positional measures’ by Lambert and Lanza (2006).

Thus:


$$y^* = \bar{y} (1 - A)^{\frac{\varepsilon-1}{\varepsilon}} \quad (19)$$

For example, consider the simple case where $\varepsilon = 0.5$ and $A = 0.3$. Substitution shows that $y^* = 1.43\bar{y}$ and a small addition to the income of anyone with an income smaller than 1.43 times the arithmetic mean produces a reduction in A . Furthermore, it can be seen that:

$$\frac{y^*}{y_\varepsilon} = (1 - A)^{-\frac{1}{\varepsilon}} \quad (20)$$

so that in this case the pivotal income is 2.04 times the equally distributed equivalent income. Further examples of the sensitivity of the pivotal income are shown in Figure 4. This shows how the ratio, y^*/\bar{y} , varies as ε is increased, given two Atkinson inequality measures of 0.6 and 0.2. Of course, for any given population, the inequality measure increases as ε is increased. Hence for a range of values of ε that are assumed to be associated with a given inequality measure, and thus different populations, the ratio y^*/\bar{y} must be higher for the lower ε values.

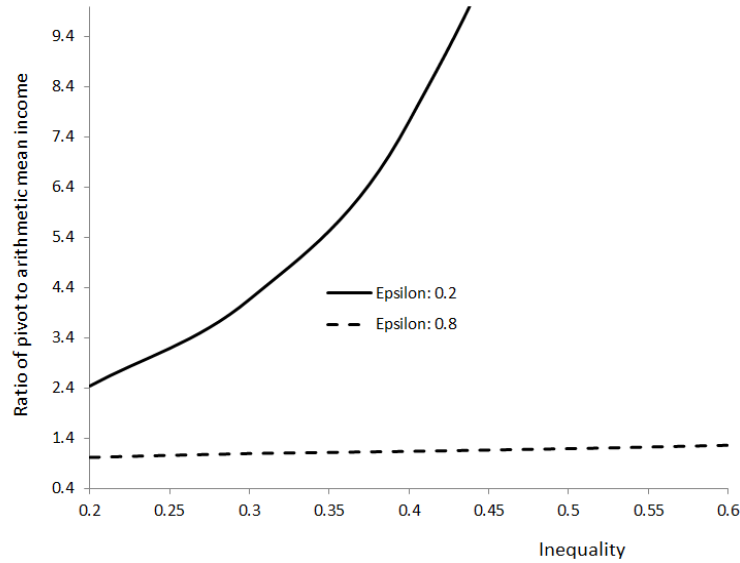
Figure 4: Ratio of Pivotal Income to Arithmetic Mean and Variation in Inequality Aversion



The variation in the ratio of the pivotal income to the arithmetic mean, as inequality increases for a given value of inequality aversion, is shown in Figure 5 for two different values of ε . It can be seen that the pivotal income is not very sensitive to variations in inequality for the higher value of inequality aversion, but is much more sensitive for lower ε .


Given the individual data used to compute the Atkinson measure, the location of the pivotal income in terms of the associated percentile can easily be obtained. However, given only the inequality measure, A , and the ratio, $y^*/\bar{y} = k$, say,

Figure 5: Variation in Pivotal Income with Inequality



along with the associated degree of inequality aversion, ε , the percentile can be determined only if the form of the income distribution is known. One approach could be to make the convenient simplifying assumption that the form of the distribution is approximated by the lognormal distribution.

Figure 6: Proportion Below Pivotal Income and Inequality: Lognormal Distribution



Suppose that y is lognormally distributed as $\Lambda(y|\mu, \sigma^2)$. In this case it can be shown that:

$$A = 1 - \exp\left(-\frac{\varepsilon\sigma^2}{2}\right) \quad (21)$$

so that, given A , the variance of logarithms is:

$$\sigma^2 = -\frac{2 \log(1 - A)}{\varepsilon} \quad (22)$$

By definition, z follows the standard Normal distribution, where:

$$z = \frac{\log(k\bar{y}) - \mu}{\sigma} \quad (23)$$

and using the fact that, for the lognormal, $\bar{y} = \exp(\mu + \sigma^2/2)$:

$$z = \frac{\log k + \sigma^2/2}{\sigma} \quad (24)$$

Figure 5 showed that there is a convex relationship between pivotal income and inequality, for high values of inequality aversion. However, Figure 6 shows that the corresponding relationship between the proportion below the pivotal income and inequality is slightly concave for the case where incomes are assumed to be lognormally distributed.

7 Conclusions

This paper has considered the problem of conveying just what it means to have a particular income inequality value, or a change in that value. One approach is based directly on the trade-off between equity and efficiency that is implied by the (abbreviated) welfare function on which the inequality measure is based. For example, given an annual growth rate, it could be said that a judge is prepared to give up a certain number of years of growth to obtain a given percentage reduction in inequality. This relies on an intuitive understanding of a sum of values, compared with a highly diverse set of values of income obtained by many individuals.

The second approach takes as its starting point the suggestion that considering the division of a cake between a small number of people provides, ‘just about the clearest idea we can have of the extent of relative inequality’, so that an inequality measure can be reduced to an equivalent expression of two income (share) values. Subramanian (2002), followed by Shorrocks (2005) with a minor modification, showed how a particularly simple summary along such lines is available for the Gini inequality measure (either in its standard form or its normalised form). This paper has concentrated on the Atkinson inequality measure, for which no simple interpretation is available: indeed, it is necessary to solve a nonlinear equation in order to calculate the income share values for any given value of Atkinson’s index.

Comparisons of ‘excess shares’ obtained for a range of assumptions, and implications for the ‘leaky bucket’ experiments, suggest that, although some interesting insights are available, it is not clear that the approach can provide the kind of information that non-specialists can easily digest in comparing different values of

inequality, at least without some discussion, including an indication of the sensitivities involved. An additional, and quite different, kind of insight is provided by the concept of the pivotal income. This was seen to be very easily calculated given the Atkinson measure (and associated degree of inequality aversion).

In view of the strong interest in inequality and changes in inequality over time, and the widespread reference to inequality measures, both in evaluating policy changes and in making a case for particular types of policy change, there is a strong incentive to try to provide empirical results that provide transparent guidance about orders of magnitude. This is especially challenging, particularly when different types of value judgement are involved and inequality measures are generally obtained as nonlinear functions of individual incomes which reduce values for many heterogeneous individuals to a single dimension. The information provided by converting inequality measures to comparisons involving a simplified cake-cutting exercise, or concentrating on a pivotal income, cannot therefore be expected to provide immediate clear answers to the communication challenge. But with careful use they may provide useful supplementary descriptions. Such approaches do at least preserve the basic value judgements behind the inequality measures and allow for the implications of alternative value judgements to be investigated. This is preferred to the use of oversimplified comparisons, and measures used for rhetorical purposes, that actually disguise the implicit value judgements of those reporting results.

Appendix A: The Gini Inequality Measure

The standard Gini inequality measure for incomes of y_1, \dots, y_n arranged in ascending order can be written as:

$$G = \frac{n+1}{n} - \frac{2}{n^2 \bar{y}} \sum_{i=1}^n (n+1-i) y_i \quad (\text{A.1})$$

If one person has income, x , and the remaining $n-1$ individuals have $(1-x)/(n-1)$ each, the arithmetic mean is $1/n$ and:

$$G = \frac{n+1}{n} - \frac{2n}{n^2} \left[x + \sum_{i=1}^{n-1} (n+1-i) \left(\frac{1-x}{n-1} \right) \right] \quad (\text{A.2})$$

$$G = \frac{n+1}{n} - \frac{2}{n} \left[x + (n+1)(1-x) - \frac{1-x}{n-1} \sum_{i=1}^{n-1} i \right] \quad (\text{A.3})$$

Using the standard result that the sum of the first n integers is $n(n+1)/2$, then $\sum_{i=1}^{n-1} i = n(n-1)/2$ and substitution and simplification gives the result that:

$$G = x - \frac{1}{n} \quad (\text{A.4})$$

Hence for any given G , the income (share) of the richer person is simply $G + \frac{1}{n}$, and G itself reflects the ‘excess share’ of the richer person. This is the result derived by Shorrocks (2005). It is immediately clear that the two-person case with $n = 2$ can be applied to provide a simplified interpretation only in the case where $G < 0.5$. Indeed, the standard Gini expression in (A.1) has a maximum value of 0.5 where there are just two individuals. Consider the distribution $[0, 1]$, for which the arithmetic mean is $1/2$. Substitution in (A.1) gives $G = \frac{3}{2} - \frac{2}{4} = \frac{1}{2}$. Unlike the Atkinson measure, the Gini has a maximum of 1 only for large n .

Shorrocks’s approach differs from the earlier analysis by Subramanian (2002), who wrote the Gini as:

$$G_s = \frac{n+1}{n-1} - \frac{2}{n(n-1)\bar{y}} \sum_{i=1}^n (n+1-i) y_i \quad (\text{A.5})$$

This expression gives the same result as the standard formulae in (A.1) for large n , but ensures that in the case of $n = 2$, G_s has a maximum of 1 for the distribution $[0, 1]$. Subramanian showed that, starting from a rank-order welfare function of the form:

$$W = \sum_{i=1}^n (n+1-i) y_i \quad (\text{A.6})$$

The normalised Gini, G_s , defined as the proportional difference between the arithmetic mean, \bar{y} , and the equally distributed equivalent income, y_{EG} , is given by:

$$y_{EG} = \frac{2}{n(n+1)} \sum_{i=1}^n (n+1-i) y_i \quad (\text{A.7})$$

This can be contrasted with the Atkinson index, a normalised measure defined in the same way but for which, as shown above, the equally distributed equivalent income is a power mean. Indeed, it can be seen that y_{EG} is a ‘reverse-order-rank-weighted’ mean of y .

Subramanian focussed on the share of income received by the poorest person, which he denoted by σ . Following a similar approach to that used above, it can be seen that:

$$\sigma = \frac{1 - G_s}{2} \quad (\text{A.8})$$

Hence, $G_s = 1 - 2\sigma$. In terms of the Shorrocks concept of excess share, $\sigma = 1 - x$, so that $G_s = 2x - 1$, which does not have the convenient interpretation obtained for G .

References

- Atkinson, A.B. (1970) On the measurement of inequality. *Journal of Public Economics*, 2, pp. 244-263
- Corvalan, A. (2014) The impact of a marginal subsidy on Gini Indices. *Review of Income and Wealth*, 60, pp. 596-603.
- Hoffmann, R. (2001) Effect of the rise of a person's income on inequality. *Brazilian Review of Econometrics*, 21, pp. 237-262.
- Lambert, P.J. (2001) *The Distribution and Redistribution of Income*. Manchester: Manchester University Press.
- Lambert, P.J. and Lanza, G. (2006) The effect on inequality of changing one or two incomes. *Journal of Economic Inequality*, 4, pp. 253-277.
- Shorrocks, A.F. (2005) *Inequality values and unequal shares*. UNU-WIDER, Helsinki.
- Subramanian, S. (2002) An elementary interpretation of the Gini inequality index. *Theory and Decision*, 52, pp. 375-379.