Title of Paper: VARIABLE POPULATIONS AND INEQUALITY-SENSITIVE ETHICAL JUDGMENTS

[Subramanian, S. (2010): 'Variable Populations and Inequality-Sensitive Ethical Judgments', in B.Basu, B. K. Chakrabarti, S. R. Chakravarti and K. Gangopadhyay (eds.): *Economics of Games, Social Choices and Quantitative Techniques*. Springer-Verlag.]

Author's Name and Institutional Affiliation:

S.Subramanian
Madras Institute of Development Studies
79 Second Main Road
Gandhinagar
Adyar
Chennai – 600 020
Tamil Nadu
INDIA.

[E-mail: **subbu@mids.ac.in**]

Acknowledgements:

This paper is based on an earlier, much shorter, version that appeared under the title 'Inequality Comparisons Across Variable Populations' in *Contemporary Issues and Ideas in Social Sciences* (2006; Volume 2, Issue 2) . The author is grateful for discussions with Satya Chakravarty, Rafael de Hoyos, Mark McGillivray, and Tony Shorrocks; and for valuable direction from Satish Jain and Sanjay Reddy. The usual caveat applies.

JEL Classification: D30, D63

Keywords: replication invariance, normalization, Gini coefficient, Atkinson index.

Word Count: 6636 words

VARIABLE POPULATIONS AND INEQUALITY-SENSITIVE ETHICAL JUDGMENTS

Abstract

This note makes the very simple point that apparently unexceptionable axioms of variable population inequality comparisons, such as the replication invariance property, can militate against other basic and intuitively plausible desiderata. This has obvious, and complicating, implications for the measurement of inequality which, for the most part, has been routinely guided by a belief in the unproblematic nature of population-neutrality principles.

VARIABLE POPULATIONS AND INEQUALITY-SENSITIVE ETHICAL JUDGMENTS

1. INTRODUCTION

Inequality comparisons are greatly facilitated when they are guided by axiom systems. (This is true also of welfare and poverty comparisons.) For the most part, the tradition has been to postulate axioms that are valid for *fixed* population comparisons. The bridge between fixed and variable population contexts has, almost entirely, been constituted by the so-called *replication invariance axiom*. Taking income, for specificity, to be the space in which inequality is appraised, replication invariance requires that how one assesses inequality should be invariant with respect to a k-fold replication of an income distribution, where k is any positive integer. The axiom, on the face of it, is unexceptionable, and is routinely treated as being innocuous: for example, Shorrocks (1988; p.433) refers to it as 'perhaps the least controversial of the "subsidiary" properties [of inequality indices]'. Recent work in poverty measurement – see Chakravarty, Kanbur and Mukherjee (2006) – however suggests that replication invariance may not be quite so un-contentious as it seems. As they put it (p.479): `Population replication axioms are now so much a part of the axiology of poverty measurement that economists take them on board without much thought.' The present article is concerned with making a similar point about the axiology of inequality comparisons.

Replication invariance is concerned with inequality comparisons which have a focus on the *proportions* of a population commanding different levels of income. However, another sort of criterion by which inequality in a society can be assessed would relate to the *absolute numbers* of the population that command, or fail to command, a preponderance of its income. Such a criterion is operationalized, in the present note, through the postulation of a pair of properties called, respectively, ‘upper pole monotonicity’ and ‘normalization’. The first of these properties requires that if all the income of a society is concentrated in the ownership of a single person, then an addition to the population of a person with identical income should result in a dilution of inequality. The second property is asymmetric in relation to the first: it requires that, given the regime of income-concentration just described, an addition to the population of a person with zero income should not be construed as worsening inequality – on the ground that, with the initial concentration of all income in a single person’s ownership, inequality is already as bad as it could possibly get. It is not hard to see that the intuition underlying a property like replication invariance could be at odds with the intuition underlying properties like upper pole monotonicity and normalization. The tension has to do with pitting considerations of relative population proportions against considerations of absolute population size – a conflict, in effect, of *fractions versus whole numbers*. The problem is elaborated on in the rest of the paper.

2. PRELIMINARY CONCEPTS

For specificity, inequality in this note will be assessed in the space of incomes. R is the set of real numbers and M is the set of positive integers. X_n is the set of all non-negative n-vectors, where n is a positive integer. A typical element of X_n is $\underline{x} = (x_1,...,x_i,...,x_n)$, where $x_i (\geq 0)$ is the income of the ith individual. Define $X \equiv \cup_{n \in M} X_n$. For every $\underline{x} \in X$, $n(\underline{x})$ is the dimensionality of $\underline{x}$, that is, $n(\underline{x}) \equiv \#N(\underline{x})$, where $N(\underline{x})$ is the set of people whose incomes are represented in $\underline{x}$. Define $X^* \equiv \{\underline{x} \in X | x_i = 0 \forall i \in N(\underline{x}) \setminus \{j\} \,\&\, \exists j \in N(\underline{x}) : x_j > 0\}$. X^*, then, is the set of income distributions which are *extremal*, in the sense of having only two types of individuals – the ‘haves’, constituted by a single individual in whose ownership the entire income of the society is concentrated, and the 'have-nots', with no income at all, who constitute the rest of the society.

Let R be a binary relation of 'inequality-sensitive' comparison defined on X. For all $\underline{x}, \underline{y} \in X$, we shall write $\underline{x} \, R \, \underline{y}$ to signify that '$\underline{x}$ reflects at most as much inequality as $\underline{y}$.' P and I are the asymmetric and symmetric parts respectively of R. For all $\underline{x}, \underline{y} \in X$, $\underline{x} \, P \, \underline{y}$ will signify that '$\underline{x}$ reflects less inequality than $\underline{y}$', and $\underline{x} \, I \, \underline{y}$ will signify that '$\underline{x}$ reflects exactly as much inequality as $\underline{y}$'. We shall take it that R is *reflexive* (for all $\underline{x} \in X : \underline{x} R \underline{x}$) and *transitive* (for all distinct $\underline{x}, \underline{y}, \underline{z} \in X : \underline{x} R \underline{y} \;\&\; \underline{y} R \underline{z} \rightarrow \underline{x} R \underline{z}$), but not necessarily complete (that is, for all distinct $\underline{x}, \underline{y} \in X$, it is not necessarily true that either $\underline{x} R \underline{y}$ or $\underline{y} R \underline{x}$ must hold). R, that is, will be taken to be a *quasi-order*. We let $\Re$ stand for the set of all quasi-orders on X.

The binary relation R will be presumed to be *anonymous*, that is, for all $\underline{x}, \underline{y} \in X$, if $\underline{y}$ is derived from $\underline{x}$ by a permutation of incomes across individuals, then $\underline{x}$ and $\underline{y}$ will be held to reflect the same extent of inequality. Inequality judgments, that is, are not influenced by personal identities. Without going into the merits of this assumption, we simply note here that the anonymity requirement is routinely invoked in the literature (for a discussion, see Broome 1989). The anonymity condition has an important part to play in a variable populations context, as exemplified in the following. Suppose $\underline{x}$ is an income distribution featuring the incomes of Hari, Rama and Krishnan, and $\underline{y}$ an income distribution featuring the incomes of Harry, Roma, Kristin and Derek. In comparing $\underline{x}$ and $\underline{y}$, it would appear that we cannot simply see $\underline{y}$ as incorporating a population addition to $\underline{x}$, and, more generally, that any 'variable population axioms' must be confined to a subset of 'time-series' comparisons and are certainly not applicable to 'cross-section' comparisons. However, the anonymity requirement finds a way out of this problem. Construct a hypothetical distribution

$\underline{z}$ in which the incomes of Harry, Roma and Kristin are identical to those of Hari, Rama and Krishnan, respectively, in $\underline{x}$. Anonymity will demand that $\underline{x}$ and $\underline{z}$ be treated identically, and $\underline{y}$ can now be seen as incorporating a population addition (in the persona of Derek) to $\underline{z}$. By invoking anonymity in this fashion, all populations of different sizes and containing different people can be compared as though one population was derived from the other through a population increment or decrement.

3. SOME AXIOMS FOR VARIABLE POPULATION INEQUALITY COMPARISONS

In what follows, we seek to impose more structure on the binary relation R by restricting it with a set of properties that may be regarded as desirable for an inequality judgment to possess. The most widely invoked restriction on inequality comparisons in a variable populations context is the property of replication invariance which – as we have seen – requires inequality judgments to be invariant with respect to population size replications. In this view, two income distributions should be treated as being identically unequal if the relative frequency distributions are identical. Formally, we have:

Replication Invariance **(*Axiom RI*)**. A binary relation $R \in \Re$ satisfies Axiom RI if and only if, for all $\underline{x}, \underline{y} \in X$ and $k \in$ M, if $\underline{y} = (\underline{x}, \ldots, \underline{x})$ and $n(\underline{y}) = kn(\underline{x})$, then $\underline{x} I \underline{y}$.

We next propose a simple criterion for inequality judgments relating to extremal income distributions. Consider an extremal distribution $\underline{x}$ of dimensionality n, such that (n-1) persons have an income of zero each and 1 person has the entire income, say x, of the society. Suppose $\underline{y}$ is derived from $\underline{x}$ through the addition of a single person with income x. For a given number (n-1) of 'have-nots' in $\underline{x}$, the number of 'haves' has risen from 1 to 2 in $\underline{y}$: this increase can naturally be associated with a dilution in the extent to which the society is polarized, and be taken to signify a reduction in inequality. This property of an inequality comparison will be called `*upper pole monotonicity*', to signify that inequality will decline monotonically with an increase in the population of the upper end of an extremal distribution:

***Upper Pole Monotonicity* (*Axiom UPM*)**. A binary relation $R \in \Re$ satisfies Axiom UPM if and only if, for all $\underline{x}, \underline{y} \in X$, if $\underline{x} \in X^*$ and $\underline{y}$ is derived from $\underline{x}$ by the addition of a single person with the same income as that of the richest person in $\underline{x}$, then $\underline{y} P \underline{x}$.

Considerations of symmetry with Axiom UPM may suggest a routine endorsement of a property such as the following one. Imagine an extremal distribution $\underline{x}$ of dimensionality n, such that (n-1) persons have an income of zero each and 1 person has a positive income of x. Suppose $\underline{v}$ is derived from $\underline{x}$ through the addition of a single person with income 0. For a given number (one, as it happens) of 'haves' in $\underline{x}$, the number of 'have-nots' has risen from (n-1) to n in $\underline{v}$: should not this increase be naturally associated with an increase in the extent to which the society is polarized, and be taken to signify an increase in inequality? A mechanical rehearsal of the reasoning underlying Axiom UPM would suggest an answer in the affirmative. However, there is a possible complication which may inhibit such a mechanical rehearsal, and this is considered in what follows.

The difference between the distributions $\underline{x}$ and $\underline{v}$ conceals a certain crucial similarity between them, which is that *both are extremal distributions*. This, indeed, is their distinctive feature. In each distribution, income is divided as unequally as it possibly could be: there is then no reason to rank the one distribution above the other in terms of inequality. In particular, it is not clear why the size of the population should enter into an assessment of the extent of inequality when, given the population size, inequality cannot get any worse. Yet, virtually all real-valued indices of inequality incorporate this irrelevant item of information. An inequality index is a mapping D:$X \to$ R, such that, for every $\underline{x} \in X$, $D(\underline{x})$ specifies a unique real number which is intended to signify the extent of inequality in $\underline{x}$. Consider, for instance, the squared coefficient of variation (C^2):

(1) For all $\underline{x} \in X$, $C^2(\underline{x}) = [1/n(\underline{x})\mu^2(\underline{x})] \sum_{i \in N(\underline{x})} x_i^2 - 1$,

where $\mu(\underline{x})$ is the mean of the distribution $\underline{x}$. If one person appropriates the entire income, the value of C^2 is (n-1). Thus, for an extremal distribution of a hundred persons, the value of C^2 is 99, while for an extremal distribution of two hundred persons, the value of C^2 is 199: it is not clear why the extent of inequality in the second case should be judged to be over twice as high as in the first case when in both cases inequality is as high as it could be. Suppose a nation starts out with an income distribution in which a single person owns all the income, and that this feature of the distribution is preserved over a period of time during which the population grows. Then it would appear to be reasonable to suggest that inequality has remained unchanged in the society, and odd to assert that inequality has increased over time. There is a piquant passage in Carroll's *Through the Looking Glass* which has relevance for this view:

> "I like the Walrus best", said Alice: "because he was a *little* sorry for the poor oysters."
>
> "He ate more than the Carpenter, though", said Tweedledee. "You see he held his handkerchief in front, so that the Carpenter couldn't count how many he took: contrariwise."

"That was mean!" Alice said indignantly. "Then I like the Carpenter best – if he didn't eat so many as the Walrus."

"But he ate as many as he could get", said Tweedledum.

This was a puzzler.

Without intending to be frivolous, one could invite the reader to think of the distribution $\underline{x}$ as the Carpenter and the distribution $\underline{y}$ as the Walrus: what we then encounter is a version of Carroll's "puzzler". Briefly, if all extremal distributions are treated as being indistinguishable in terms of the extent of inequality, then this would be a case against postulating a property – call it 'Lower Pole Monotonicity' - which is derived as a mirror image of 'Upper Pole Monotonicity'. Rather, the case would be in favour of a sort of 'weak normalization' which asserts that additions of zero-income individuals to an extremal distribution should not be construed as worsening inequality[1]:

***Weak Normalization* (*Axiom WN*)**. A binary relation $R \in \Re$ satisfies Axiom WN if and only if, for all $\underline{x}, \underline{y} \in X$, if $\underline{x} \in X^*$ and $\underline{y}$ is derived from $\underline{x}$ by the addition of a single person with zero income, then $\neg(\underline{x} P \underline{y})$.

Axiom WN demands only that an extremal distribution of smaller population should not be declared to be inequality-wise *preferred* to an extremal distribution of larger population. A stronger condition – call it 'normalization' – might call for declaring the two distributions to be inequality-wise *indifferent*:

***Normalization* (*Axiom N*).** A binary relation $R \in \Re$ satisfies Axiom N if and only if, for all $\underline{x}, \underline{y} \in X$, if $\underline{x} \in X^*$ and $\underline{y}$ is derived from $\underline{x}$ by the addition of a single person with zero income, then $\underline{x} \, I \, \underline{y}$.

The normalization axiom is not a property of any commonly employed real-valued measure of inequality. For an extremal distribution, the value of the Gini inequality coefficient is $(n-1)/(n+1)$, which is an increasing function of n; for the squared coefficient of variation (as we have seen), it is $(n-1)$, again an increasing function of n; for one of the two Theil indices, it is log n, also an increasing function of n. It appears completely arbitrary that when a distribution is extremal, it should be differentially penalized the larger its dimension is. For future reference, we present a normalized version of the squared coefficient of variation, C^2*.

(2) For all $\underline{x} \in X : C^2*(\underline{x}) = [1/(n(\underline{x})-1)][(1/n(\underline{x})\mu^2(\underline{x})) \sum_{i \in N(\underline{x})} x_i^2 - 1]$ $[= \{1/(n(\underline{x})-1)\}C^2(\underline{x})]$.

In what follows, we examine the sorts of inequality judgments that are possible when they are required to satisfy certain combinations of the axioms we have discussed.

4. ON THE POSSIBILITY OF CONSISTENT INEQUALITY COMPARISONS

The following proposition is true.

Proposition. (i) There exists a binary relation $R \in \Re$ which satisfies Upper Pole Monotonicity and Normalization; and (ii) there exists a binary relation $R \in \Re$ which satisfies Upper Pole

Monotonicity and Replication Invariance; but (iii) there exists no binary relation $R \in \Re$ which satisfies Replication Invariance, Upper Pole Monotonicity and Weak Normalization.

Proof. (i) We claim that the requirements stated in part (i) of the Proposition are satisfied by the binary relation *R**, defined as follows:

$\forall \underline{x}, \underline{y} \in X$, $\underline{x} R^* \underline{y}$ if and only if $C^2 *(\underline{x}) \leq C^2 *(\underline{y})$, where the index $C^2 *$ is as defined in equation (2).

Note first that *R** is obviously reflexive. To check for transitivity, suppose there exists a triple $\{\underline{x}, \underline{y}, \underline{z}\}$ which is a subset of *X*, such that $\underline{x} R^* \underline{y}$ and $\underline{y} R^* \underline{z}$: this is predicated on the truth of $C^2 *(\underline{x}) \leq C^2 *(\underline{y})$ and $C^2 *(\underline{y}) \leq C^2 *(\underline{z})$, and since the relation ≤ defined on real numbers is transitive, we must have $C^2 *(\underline{x}) \leq C^2 *(\underline{z})$ which, by construction of *R**, implies $\underline{x} R^* \underline{z}$, and this proves the transitivity of *R**. Next, to check on Axiom UPM, consider any $\underline{x} \in X^*$ of dimensionality *n*, and let $\underline{y}$ be derived from $\underline{x}$ through the addition of a single person with the same income, say *x*, as that of the richest person in $\underline{x}$. We need to show that $\underline{y} P^* \underline{x}$. Given (2) it is easy to check that $C^2 *(\underline{y}) = (n-1)/2n < C^2 *(\underline{x}) = 1$, that is, $\underline{y} P^* \underline{x}$, as needed for *R** to satisfy Axiom UPM. This leaves us with Axiom N to verify. Consider any $\underline{x}, \underline{y} \in X^*$ such that the number of persons with zero income in $\underline{x}$ is *p* and the number of persons with zero income in $\underline{y}$ is (*p*+1), while the richest person in both distributions has the same income. We need to show that $\underline{x} I^* \underline{y}$. Given (2) it can be verified that, since $C^2 *(\underline{x}) = C^2 *(\underline{y}) = 1$, it follows, by construction of *R**, that $\underline{x} I^* \underline{y}$, that is, *R** satisfies Axiom N.

(ii) We claim that the requirements stated in part (ii) of the Proposition are satisfied by the binary relation $\hat{R}$, defined as follows:

$\forall \underline{x}, \underline{y} \in X$, $\underline{x}\ \hat{R}\ \underline{y}$ if and only if $C^2(\underline{x}) \leq C^2(\underline{y})$, where the index C^2 is as defined in equation (1).

Since the claim can be easily substantiated, and along the same lines as the demonstration in part (i) of the Proposition, a detailed proof is omitted.

(iii) A simple counter-example suffices to prove part (iii) of the Proposition. Let *x* be any positive real number, and consider the following distributions $\underline{x}$, $\underline{y}$ and $\underline{z}$ belonging to *X*: $\underline{x}$ = (0,0,*x*,*x*), $\underline{y}$ = (0,0,*x*), and $\underline{z}$ = (0,*x*). Then:

By the upper pole monotonicity axiom,

(P1) $\underline{x} P \underline{y}$.

By the axiom of replication invariance,

(P2) $\underline{z} I \underline{x}$.

By virtue of transitivity of *R*, and given (P1) and (P2):

(P3) $\underline{z} P \underline{y}$.

However, by the weak normalization axiom,

(P4) $\neg(\underline{z} P \underline{y})$.

From (P3) and (P4) we have a contradiction. This completes the proof of the Proposition. ■

The proof of part (i) of the Proposition above revolves around the construction of an inequality index – C^2* - which must be judged to be a peculiar index in terms of common convention: it violates replication invariance, which is not a feature of any commonly employed inequality index. The proof of part (ii), however, revolves around an inequality index - C^2 - which violates normalization, and this is a feature of all commonly employed inequality indices. The point about replication invariance and normalization is at the heart of the (impossibility) result contained in part (iii) of the Proposition. This leads naturally to a consideration of the real-valued representation of inequality under alternative specifications of the axiom system by which the aggregation procedure is constrained. In particular, it is of interest to examine some consequences of measuring inequality when the inequality index is required to satisfy (a) replication invariance at the expense of normalization, and (b) normalization at the expense of replication invariance. One aspect of this problem is discussed in what follows.

5. AGGREGATION: NORMALIZATION VERSUS REPLICATION INVARIANCE?

When inequality is measured in terms of a real-valued index, the conflict between replication invariance and normalization is reflected sharply in one particular interpretation of the inequality index. This interpretation revolves around establishing a correspondence between the value of an inequality measure for an n-person distribution and the shares in which a cake of given size is split between two persons. If such an equivalence can be demonstrated, then this would be a very useful outcome, because, in many ways, our intuitive grasp of inequality is clearest and sharpest in the context of a two-person cake-sharing exercise. As it happens, the correspondence in question can, indeed, be effected – with or without qualification. Two alternative approaches to the problem are available in Shorrocks (2005) and Subramanian (1995, 2002). The difference in the results obtained by the two authors resides in the fact that Shorrocks considers inequality measures which satisfy replication invariance, while Subramanian considers inequality measures which satisfy normalization. As discussed below, normalization at the expense of replication invariance affords a rather more unambiguous link between n-person and two-person distributions than does replication invariance at the cost of normalization.

Shorrocks (2005) demonstrates that the Gini coefficient (which, of course, is a replication-invariant inequality measure) can be interpreted as the `excess share' of the richer of two persons when a cake of given size is split between two individuals. The `fair share' in a two-person situation is one-half; and a Gini coefficient (G) of 0.4, as it turns out, can be interpreted as the share in excess of 0.5 going to the richer person: in this interpretation, a Gini coefficient of 0.4 for an n-person distribution is equivalent to the richer of two persons, in a two-way division of a cake, receiving 90 per cent of the cake. When, however, G for an n-person distribution exceeds 0.5, the corresponding share of the poorer of the two persons in an 'equivalent' cake-sharing setting would have to be negative – and negative shares are not easy to get an intuitive handle on. Shorrocks shows that the `excess share' interpretation is valid not only for a two-person split but for a general, n-person split, in terms of which, G is the excess of the richest person's share, call it r, over his fair share, which is just $1/n$, so that $r = G + 1/n$. Shorrocks uses the term 'modulo 2' to indicate excess shares in the context of a 2-person split, and the term 'modulo 10' to indicate excess shares in the context of a 10-person split. Thus, while a G-value of , say, 0.7 would be equivalent to a hard-to-interpret 120 per cent share for the richer person in a cake split two ways, it would also be equivalent to a readily comprehensible share of 80 per cent for the richest person in a cake split 10 ways, with the poorest 9 individuals sharing the balance 20 per cent equally among themselves. Notice that G can be as high as 0.9 before r begins to exceed 100 per cent of the cake in a 10-way split. Values of G in excess of 0.9 are not generally encountered in actual empirical distributions, and therefore, in practice, the 'modulo 10' interpretation of Gini should not pose the problem of having to interpret negative shares or shares in excess of unity.

Shorrocks indicates that the 'excess share' interpretation can be applied to a range of inequality measures, including the class of Generalized Entropy measures and the family of 'ethical' indices due to Atkinson (1970). As he puts it (Shorrocks 2005; p.4): 'If one …

considers a distribution consisting of one rich person with the income share r and (n-1) poorer people each with income share (1- r)/(n-1), the values of each of these indices may be written as increasing functions of the excess share of the richest person, $r - 1/n$. However, the relationship [between] the inequality value and the excess share is more complex and, as a consequence, the interpretation is less immediate.' (I have taken some minor notational liberties in reproducing the quote.) While the 'excess share' interpretation of an inequality index is of very considerable interpretive value, the keen edge of immediate intuitive clarity does get blunted when one departs from the modulo 2 interpretation. Subramanian (1995, 2002) indicates that if an inequality index is required to satisfy normalization at the expense of replication invariance, then it is possible to preserve the '2-person split' interpretation of the inequality value, without risking shares in excess of 100 per cent for the richer person and negative shares for the poorer person. This can be shown to hold for normalized versions of the Gini coefficient (Subramanian 2002) and the Atkinson class of indices (Subramanian 1995). The relevant results are briefly reviewed in what follows.

It may be recalled that Atkinson (1970) sought to relate the extent of inequality in any income distribution to the loss in welfare caused by the presence of inequality. To operationalize this approach requires, first, the specification of some appropriate ('equity-sensitive') welfare function defined on an income distribution. Thus, given any income vector $\underline{x}$, with mean $\mu(\underline{x})$, let $W(\underline{x})$ be the welfare level associated with the given distribution of incomes. The *equally distributed equivalent income* is that level of income, call it x^e, such that its equal distribution leads to a level of welfare which is the same as the welfare level associated with the distribution $\underline{x}$ under review. Given any $\underline{x}$, and a welfare function W defined on it, a measure of inequality D for the distribution can be obtained as the proportionate difference between the mean of the distribution and the equally distributed equivalent income:

(3) $D(\underline{x}) = [\mu(\underline{x}) - x^e(\underline{x})] / \mu(\underline{x})$.

The Gini coefficient of inequality can be derived from a 'Borda' welfare function (see Sen, 1973), in terms of which aggregate welfare is a rank-order weighted sum of individual incomes. If $\underline{x}$ is a non-decreasingly ordered vector of incomes, then the Borda social welfare function is given by

(4) $W^B(\underline{x}) = \sum_{i \in N(\underline{x})} (n+1-i)x_i$.

The equally distributed equivalent income for the Borda welfare function – call it x_B^e - is given by:

(5) $x_B^e(\underline{x}) = [2/n(n+1)] \sum_{i \in N(\underline{x})} (n+1-i)x_i$.

The Gini coefficient *G* is the proportionate difference between $\mu(\underline{x})$ and $x_B^e(\underline{x})$. Given (3) and (4), we have:

(6) $G(\underline{x}) = 1 - [2/n(n+1)\mu] \sum_{i \in N(\underline{x})} (n+1-i)x_i$.

A normalized version of *G*, call it *G**, can be obtained as the ratio of the difference between μ and x_B^e and the difference between μ and the lowest value x_{B0}^e which x_B^e can attain:

(7) $G^*(\underline{x}) = [\mu(\underline{x}) - x^e(\underline{x})]/[\mu(\underline{x}) - x^e_{B0}(\underline{x})]$.

In turn, x^e_B attains its lowest value when the distribution is maximally concentrated (with the richest person appropriating the entire income), and in such a situation, it can be verified that

(8) $x^e_{B0}(\underline{x}) = 2\mu(\underline{x})/(n+1)$.

Given (7) and (8), we can now write the normalized version of the Gini coefficient as:

(9) $G^*(\underline{x}) = \frac{n+1}{n-1} - \frac{2}{n(n-1)\mu} \sum_{i \in N(\underline{x})} (n+1-i)x_i$.

From (6) and (9), we obtain:

(10) $G^*(\underline{x}) = \left[\frac{n+1}{n-1}\right] G(\underline{x})$.

Notice from (6) that the highest value *G* can attain is (*n*-1)/(*n*+1), which approaches unity as *n* becomes larger, whereas the highest value *G** can attain is always exactly unity, irrespective of the dimensionality of the income vector. *G**, unlike *G*, is normalized but not replication-invariant. This fact enables one to obtain a very straightforward result for *G** on the relationship between the value of *G** and the share of the poorer of two individuals when a cake is split two ways. To see what is involved, given any ordered *n*-vector of incomes $\underline{x}$, define a *dichotomously allocated equivalent distribution* (*daed*) as a non-decreasingly ordered 2-vector $\underline{x}^* \equiv (x_1^*, x_2^*)$ such that $\underline{x}^*$ has the same mean $\mu(\underline{x})$ as $\underline{x}$ and the same normalized Gini $G^*(\underline{x})$ as $\underline{x}$. One then obtains a pair of simultaneous equations in x_1^* and

x_2^*; solving for these, and letting σ_G stand for the income share $x_1^* / 2\mu(\underline{x})$ of the poorer individual in the *daed* $\underline{x}^*$, it can be verified (see Subramanian 2002 for details) that

(11) $\sigma_G = (1 - G^*)/2$.

Thus, if G^* for some n-person distribution should be 0.4, this is equivalent to a situation in which the poorer of two persons gets a 30 per cent share of a cake that is split two ways. The significance of the normalized Gini can always be understood in terms of this helpful `2-person split'. The Atkinson class of inequality indices can be similarly interpreted, as is discussed below in the light of Subramanian (1995).

Atkinson employs a utilitarian social welfare function which is a sum of identical individual utility functions that are symmetric, increasing and strictly concave, and specialized to the 'constant elasticity-of-marginal utility' form. The utility function, for each person i, is given by

(12) $u(x_i) = \frac{1}{\lambda} x_i^{\lambda}, \lambda \in (0,1)$.[2]

Given any income vector $\underline{x}$, the Atkinson welfare function defined on $\underline{x}$ can be written as:

(13) $W(\underline{x})[= \sum_{i \in N(\underline{x})} u(x_i)] = (1/\lambda) \sum_{i \in N(\underline{x})} x_i^{\lambda}, \lambda \in (0,1)$.

It is easy to check that the equally-distributed equivalent income for the Atkinson social evaluation function is

(14) $x_A^e(\underline{x}) = [\frac{1}{n}\sum_{i\in N(\underline{x})} x_i^\lambda]^{1/\lambda}, \lambda \in (0,1).$

Atkinson's inequality index is then given by:

(15) $A(\underline{x})[=\{\mu(\underline{x}) - x_A^e(\underline{x})\}/\mu(\underline{x})] = 1 - [\frac{1}{n\mu^\lambda}\sum_{i\in N(\underline{x})} x_i^\lambda]^{1/\lambda}, \lambda \in (0,1)$.

If x_{A0}^e is the minimum value which x_A^e can attain (corresponding to a situation in which all income is concentrated in a single person's hands), then it is easy to check that

(16) $x_{A0}^e(\underline{x}) = n^{\frac{\lambda-1}{\lambda}}\mu(\underline{x}), \lambda \in (0,1)$.

A normalized version of the Atkinson inequality index A can be obtained as the ratio of the difference between μ and x_A^e and the difference between μ and x_{A0}^e: this index – call it A^* - always attains a value of unity, irrespective of the dimensionality of the distribution, when the latter is an extremal one. Given (14) and (16), it can be verified that

(17) $A^*(\underline{x}) = \left[\frac{1}{1-n^{(\lambda-1)/\lambda}}\right]\left[1 - (\frac{1}{\mu(\underline{x})})(\frac{1}{n}\sum_{i\in N(\underline{x})} x_i^\lambda)^{1/\lambda}\right], \lambda \in (0,1)$.

In view of (15) and (17), the relationship between A^* and A is defined by:

(17) $A^*(\underline{x}) = \left[\frac{1}{1-n^{(\lambda-1)/\lambda}}\right]A(\underline{x}), \lambda \in (0,1)$.

Providing a simple interpretation for the index A*, to recall, is the motivation with which we started out. Given an *n*-person distribution $\underline{x}$ and some value for the 'inequality-aversion' parameter λ, it is not the easiest of things to conceptualize precisely what a particular value of the inequality measure $A^*(\underline{x})$ 'really means' in terms of categories of inequality that we may be familiar with at a more 'primitive' level. The notion of a 'dichotomously allocated equivalent distribution (or *daed*)', as we have seen, is of help in this context. Given $\underline{x}$, the corresponding *daed* is the ordered 2-vector $\underline{x}^* = (x_1^*, x_2^*)$ such that $\underline{x}^*$ and $\underline{x}$ have both the same means and the same values of the (normalized) Atkinson inequality index. Then, $\mu(\underline{x}^*) = \mu(\underline{x})$ implies

(18) $x_1^* + x_2^* = 2\mu(\underline{x})$;

and $A^*(\underline{x}^*) = A^*(\underline{x})$ implies – making appropriate use of (17) – that

(19) $$\frac{\mu(\underline{x}) - \left[\frac{1}{2}\{(x_1^*)^\lambda + (x_2^*)^\lambda\}\right]^{1/\lambda}}{\left[1 - 2^{(\lambda-1)/\lambda}\right]\mu(\underline{x})} = A^*(\underline{x}) \ .$$

Substituting for x_2^* from (18) into (19), and after suitable manipulation of (19), we obtain:

(20) $(x_1^*)^\lambda + (2\mu(\underline{x}) - x_1^*)^\lambda = 2[\mu(\underline{x})\{1 - A^*(\underline{x})(1 - 2^{(\lambda-1)/\lambda})\}]^{1/\lambda}$.

Designating the poorer person's income share in the distribution $\underline{x}^*$ by σ_A, we have: $\sigma_A = x_1^* / 2\mu(\underline{x})$, whence $x_1^* = 2\mu(\underline{x})\sigma_A$; substituting for x_1^* into (20) and simplifying, yields:

(21) $\sigma_A^\lambda + (1-\sigma_A)^\lambda = 2^{1-\lambda}[1 - A^*(\underline{x})(1-2^{(\lambda-1)/\lambda})]^\lambda$.

Using (21), we can solve for σ_A, given any value of *A** (though a closed-form solution expressing σ_A as a function of *A** is not available). Notice from (21) that when $A^*(\underline{x}) = 0$ (no inequality), $\sigma_A = 1/2$ (equal share), and when $A^*(\underline{x}) = 1$ (perfect concentration), $\sigma_A = 0$ (the poorer person receives nothing). In general, given the value of the index *A** for any *n*-person distribution, one can transform it into an 'equivalent' value of σ_A - the share of the poorer person in a 2-person distribution – which affords an immediate and vivid picture of the extent of inequality that *A** 'signifies'. Importing the normalization axiom at the expense of the replication-invariance axiom into the aggregation exercise facilitates this clear and unqualified equivalence result.

A brief digression is in order here. In order to compute the values of σ_G and σ_A, we need to know the values of *G** and *A**. However, when we work with the estimates of inequality computed by other scholars, typically what would be available to us are the values of *G* and *A*, not of *G** and *A** (because, conventionally, the Gini coefficient and the Atkinson index are presented in their replication-invariant form). If we knew the value of *n*, we could, of course, compute the values of *G** and *A** from the corresponding values of *G* and *A*, by employing equations (10) and (17) respectively. But then, it is unlikely that we would have access to knowledge concerning the size of *n*, though we may be reasonably confident that, for instance, *n* is at least 200 but not larger than 200,000. Under these circumstances, and given (10) and (17), it is of interest to ask under what condition *G** (respectively, *A**) can be deemed to be a ‘valid’ approximation of *G* (respectively, *A*). The answer would depend on the extent of our tolerance of the discrepancy between the normalized and the non-normalized versions of the relevant inequality index. If *D** is the normalized, and *D* the non-normalized, value of some inequality index that we are interested in, let ε_D be a measure of the proportionate distance between *D* and *D**, where ε_D is defined as: $\varepsilon_D \equiv 1-\min(D,D^*)/\max(D,D^*)$. Given (10), it is easy to see that, for the Gini coefficient, $\varepsilon_G = 2/(n+1)$, while for the Atkinson index, it can be verified from (17), that $\varepsilon_A = n^{(\lambda-1)/\lambda}$. Let ε^* be the value of ε that we are prepared to tolerate: that is, if ε_G (respectively, ε_A) is not greater than ε^*, then we shall say that *G** (respectively, *A**) is a ‘valid’ approximation of *G* (respectively, *A*). Given the expressions for ε_G and ε_A, it is easily verified that *G** (respectively, *A**) can be accepted as a ‘valid’ approximation of *G* (respectively, *A*) if $n \geq 2/\varepsilon^* - 1$ (respectively, if $n \geq (\varepsilon^*)^{\lambda/(\lambda-1)}$). If ε^* is set at 1 per cent (so that, for instance, we are prepared to accept 0.396 as a valid approximation of 0.4), then it turns out that *G** is an acceptable approximation of *G* provided *n* is at least 199; for the Atkinson index, if λ is 0.5, *A** is `acceptably close’ to *A* whenever *n* is at least 100; however, if λ is 0.9, the test of acceptability will be passed only if *n* is at least 10^{18}. With the knowledge that *n* lies between 200 and 200,000, we can employ *G* as an acceptable surrogate for *G**, and *A* as an acceptable surrogate for *A** provided λ does not exceed 0.723.

The immediately preceding considerations are relevant for the small empirical illustration presented in Table 1. In this Table, the values of σ_A and σ_G are provided for corresponding values of *A** and *G**, derived from the values of *A* and *G* respectively, that have been computed from 1970 income distribution data for Malaysia by Anand (1983). We can see from the Table that, for example, when $\lambda = 0.5$, the welfare considerations underlying the (normalized) Atkinson index imply that the extent of interpersonal inequality that obtains in the distribution of income in Malaysia is ‘comparable’ to a situation in which the poorer of two persons receives just a little under a fifth of a cake that has to be divided two ways (this share is just a little over a quarter for the normalized Gini coefficient): this ‘equivalence translation’ affords a graphic way of comprehending inequality.

[Table 1 to be inserted here]

6. CONCLUDING OBSERVATIONS

Derek Parfit (1984) has alerted us to the serious possibility that variable population situations can present a major challenge to one’s moral intuition. This paper is a specific, and very simple, example of this proposition, as applied to the exercise of effecting inequality-sensitive ethical comparisons. In assessing the import of the Proposition presented in Section 4, it seems to be hard to quarrel with the Upper Pole Monotonicity Axiom. The real problem is the conflict between Weak Normalization and Replication Invariance. An attempt has been made in this note to rationalize Weak Normalization. If there is something to be said for the rationale provided, then this should cast doubt on the readiness with which Replication Invariance has been routinely accepted in much of the literature on inequality measurement. This also has implications for the aggregation problem in inequality measurement – the problem of constructing ‘satisfactory’ real-valued inequality indices which must be informed by a deliberate judgment on whether to sacrifice Replication Invariance in the cause of Normalization or the other way around. The compulsion for conscious choice is precipitated by the fact that individually attractive principles of inequality comparison may not always be collectively coherent.

Table 1: Alternative Inequality Indices and the Corresponding 'Two-Way Splits' for the Distribution of Individuals by Per Capita Household Income: Malaysia 1970.

λ	A	A^*	σ_A (per cent)	G	G^*	σ_G (per cent)
0.75	0.1162	0.1185	28.34	-	-	-
0.50	0.2124	0.2124	19.19	-	-	-
0.25	0.3026	0.3026	17.40	-	-	-
0.10	0.3807	0.3807	9.32	-	-	-
0.01	0.7615	0.7615	1.37	-	-	-
-	-	-	-	0.4980	0.4980	25.10

Note:

(1) Table 1 is based on data furnished in Tables 3-8 and 3-9 of Anand (1983).

(2) Anand's computations are based on data in the Malaysian Post-Enumeration Survey of 1970: the sample size of the survey (number of individuals) is reported to be 134,186.

(3) In Table 1, λ stands for the inequality aversion parameter in the Atkinson inequality index; A stands for the Atkinson inequality measure; A^* stands for the normalized Atkinson inequality measure; G stands for the Gini coefficient of inequality; G^* stands for the normalized Gini; and σ_A (respectively, σ_G) stands for the 'equivalent' income share, corresponding to the normalized Atkinson measure (respectively, normalized Gini coefficient) of the poorer of two persons in the dichotomously allocated equivalent distribution [see equations (21) and (11) in the text].

(4) Anand's reported values of A (Table 3-9 in Anand, 1983) have been mapped into their corresponding values of A^* by using the relation specified in equation (17) in the text:

$A^*(\underline{x}) = \left[\frac{1}{1-n^{(\lambda-1)/\lambda}}\right] A(\underline{x}), \lambda \in (0,1)$; and Anand's reported value of *G* (Table 3-8 in Anand, 1983) has been mapped into its corresponding value of *G** by using the relation specified in equation (10): $G^*(\underline{x}) = \left[\frac{n+1}{n-1}\right] G(\underline{x})$. *n*, to recall, is 134,186.

(5) Notice that there is no discrepancy between the values of *G** and *G*, and – except in the case of $\lambda = 0.75$ – between the values of *A** and *A*, up to the fourth decimal place.

Source:

Based on computed values of *G* and *A* furnished in Tables 3-8 and 3-9 respectively of Anand (1983).

NOTES

1. There is a second possible reason why 'additions of zero-income individuals to an extremal distribution should not be construed as worsening inequality'. This reason is not crucial to the concerns of this note, but is dealt with here for purposes of completeness, and also because it echoes a concern voiced in Kanbur and Mukherjee (2006). It is worth noticing, in the context of the upper pole monotonicity axiom, that the transition from $\underline{x}$ to $\underline{y}$ has been mediated by the addition of a human life to the society under review: the judgment that $\underline{y}$ is to be inequality-wise preferred to $\underline{x}$ upholds also the value of the existence of a certain type of person (a 'have' person). While the transition from $\underline{x}$ to $\underline{v}$ is similarly mediated by the addition of a human life to the society under review, the judgment that $\underline{x}$ is to be inequality-wise preferred to $\underline{v}$ is now unhappily accompanied by a *dis*-valuing of the existence of a certain type of person (a 'have-not' person). If we wish to avoid inequality judgments that go against the existence of certain categories of lives, then we should not subscribe to the notion that $\underline{x}$ is inequality-wise preferred to $\underline{v}$. The force of such an argument is strong if the statement '$\underline{a}$ is inequality-wise better than $\underline{b}$' is always interpreted to mean also that '$\underline{a}$ is all-things-considered ethically better than $\underline{b}$'. Does this entailment necessarily hold? One could contend not. It can be maintained that the betterness relation *R* employed here refers only to inequality judgments, not to overall ethical judgments: in this view, $\underline{a}$ can be judged to be inequality-wise better than $\underline{b}$, though, on a balance of other ethical considerations (such as ones relating to the value of life), $\underline{b}$ can be judged to be overall a better state of affairs than $\underline{a}$. This is a coherent position to adopt - although, on the other hand, one could also ask: if $\underline{a}$ is inequality-

wise better than $\underline{b}$ without entailing that $\underline{a}$ is better than $\underline{b}$, then what (to echo Broome 1989) `is the good of equality'? The arguments on either side may have to be treated as inconclusive. For those who cannot endorse life-disvaluing inequality judgments, there is sufficient reason to refuse to rank the distribution $\underline{x}$ over the distribution $\underline{y}$. For those who would draw a distinction between inequality comparisons and more comprehensive ethical comparisons into which inequality judgments may be factored and given more or less weight, some other line of reasoning may be required to persuade them against pronouncing '$\underline{x}P\underline{y}$'. It is this alternative line of reasoning that has been explored in the text.

2. Non-positive values of λ are not considered, because of the problems occasioned by these in the presence of zero-incomes in a distribution (for a discussion of which see Anand 1983; pp. 84-86).

REFERENCES

Anand, S. (1983). *Inequality and Poverty in Malaysia: Measurement and Decomposition*. New York: Oxford University Press.

Atkinson, A.B. (1970). `On the Measurement of Inequality', *Journal of Economic Theory*, 2, 244-263.

Broome, J. 1989. 'What's The Good Of Equality?' In J. D. Hey (ed.) *Current Issues in Microeconomics*, London: Palgrave Macmillan.

Chakravarty, S., S. R. Kanbur and D. Mukherjee. 2006. 'Population Growth and Poverty Measurement', *Social Choice and Welfare* 26(3): 471-483.

Kanbur, R. and D. Mukherjee (2006). 'Premature Mortality and Poverty Measurement'. (October 2006). Available at http://www.arts.cornell.edu/poverty/kanbur/Pov&Death.pdf

Parfit, D. 1984. *Reasons and Persons*, Oxford: Clarendon Press.

Sen,A.K. (1973). *On Economic Inequality*. Oxford: Clarendon Press.

Shorrocks, A. F. (1988). 'Aggregation Issues in Inequality Measurement', In W. Eichhorn (ed.) *Measurement in Economics: Theory and Applications in Economic Indices*, Heidelberg: Physica-Verlag.

Shorrocks, A. F. (2005). 'Inequality Values and Unequal Shares', UNU-WIDER, Helsinki. Available at http://www.wider.unu.edu/conference/conference-2005-5/conference-2005.5.htm

Subramanian, S. (1995). 'Toward a Simple Interpretation of the Atkinson Class of Inequality Indices', MIDS Working Paper No. 133, Madras Institute of Development Studies: Chennai. (*mimeo*.)

Subramanian, S. (2002). 'An Elementary Interpretation of the Gini Inequality Index', *Theory and Decision*, 52(4): 375-379.