

Research Paper No. 2005/59

Reckoning Inter-group Poverty Differentials in the Measurement of Aggregate Poverty

S. Subramanian *

September 2005

Abstract

In a heterogeneous population which can be partitioned into well-defined subgroups, it is plausible that the extent of measured aggregate poverty should depend upon the distribution of poverty across the subgroups. A judgment in favour of an *equal* inter-group distribution of poverty could arise in two ways. In the first approach, equality is upheld as an *intrinsic* social virtue, and the aggregate measure of poverty, in line with this view, is ‘adjusted’ to reflect the extent of inter-group disparity in the distribution of poverty that obtains. In the present paper, this approach is examined, with specific reference to the advancement of a diagrammatic aid to analysis called the *group poverty profile*. In the second approach, equality is upheld for *instrumental* reasons which arise from the observed fact that any individual’s level of deprivation is a function not only of one’s own income, but of the general level of prosperity of the group to which one is affiliated. Individual deprivation functions are specialized to a form which reflects this ‘group-affiliation’ externality, and the resulting poverty measure is studied with respect to its properties, and its implications for inter-group equity. The analysis is briefly extended to a review of the measurement of literacy, along externality-motivated lines suggested elsewhere by Basu and Foster. The paper concludes that social realism in the measurement of deprivation is often compromised by mainstream approaches to economic theorizing in which both heterogeneity and group-related externalities are generally de-emphasized.

Keywords: group poverty profile, group Lorenz profile, externality, group affiliation

JEL classification: D63, I32

Copyright © UNU-WIDER 2005

* UNU-WIDER, Helsinki, and Madras Institute of Development Studies, Chennai; subbu@mids.ac.in

This study has been prepared within the 2005 UNU-WIDER Sabbatical and Visiting Fellows programme, and is published in the project on New Directions in Development Economics.

UNU-WIDER acknowledges the financial contributions to the research programme by the governments of Denmark (Royal Ministry of Foreign Affairs), Finland (Ministry for Foreign Affairs), Norway (Royal Ministry of Foreign Affairs), Sweden (Swedish International Development Cooperation Agency—Sida) and the United Kingdom (Department for International Development).

ISSN 1810-2611 ISBN 92-9190-742-1 (internet version)

Acknowledgements

The author is grateful to Professor Prasanta Pattanaik for a very helpful discussion of the subject of this paper, to D. Jayaraj for earlier collaborative work on related themes, and to Subroto Mukherjee for comments. The usual *caveat* applies.

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

www.wider.unu.edu

publications@wider.unu.edu

UNU World Institute for Development Economics Research (UNU-WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Camera-ready typescript prepared by Adam Swallow at UNU-WIDER

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

1 Introduction

Two aspects of social reality frequently encountered are that populations are heterogeneous and that deprivation at an individual level is a function of the deprivation status of the group to which the individual is affiliated. The first aspect could pave the way for measuring aggregate poverty in such a way that group-related disparities in the distribution of poverty are explicitly taken into account and penalized. The second could pave the way for measuring poverty in such a way that individual deprivation is seen to depend on group deprivation, so that ‘group-affiliation externalities’ are explicitly reckoned in the assessment of overall poverty. In either case, it turns out that inter-group equality in the distribution of poverty is a valued outcome—for intrinsic reasons of the desirability of equality as a social virtue in the first case, and for instrumental reasons arising from the link between externality and equity in the second case. In the present paper, both the ‘intrinsic’ and the ‘instrumental’ cases for reckoning inter-group differentials in the measurement of aggregate poverty are investigated.

The ‘intrinsic’ approach has been examined by Anand and Sen (1995) in the context of deriving a ‘gender-adjusted human development index’, and variations on this theme have been explored by Jayaraj and Subramanian (1999), Majumdar and Subramanian (2000), Subramanian and Majumdar (2002), and Subramanian (2004b). The Subramanian and Majumdar paper provides an axiomatic justification for a specific ‘group inequality-adjusted’ aggregate index of deprivation. In this paper, a diagrammatic link to that index, in the form of a graph called the ‘group poverty profile’, which is directly inspired by Shorrocks’ (1995, 1996) ‘poverty gap profile’ employed in a different context, is advanced. The poverty gap profile could prove to be a useful instrument (not unlike the Lorenz curve in inequality measurement) for comparing alternative regimes of the inter-group distribution of poverty.

The ‘instrumental’ approach has been explored by Basu and Foster (1998) in the context of deriving an ‘externality-adjusted’ measure of literacy. An application to the poverty context can also be found in Jayaraj and Subramanian (2000). The underlying idea here is that deprivation at the level of an individual is a function not only of her own income but also of the average income of the group to which she belongs. A poverty index which displays this sort of group-related externality is presented as an illustrative outcome of this approach to measuring aggregate poverty, and its implications for inter-group equality in the distribution of poverty, as well as for poverty rankings and poverty axiomatics, are examined. It is also shown—through a consideration of the literacy measurement problem—that the ‘externality’ approach is a general one with useful specific applications to the measurement of various social indicators.

Both approaches to reckoning subgroup poverty in the measurement of overall poverty suggest that departures from the conventional assumptions of ‘homogeneity’ and ‘individualism’ can have non-trivial implications for how we view and measure poverty. In addressing these issues, this paper has been organized as follows. Section 2 deals with the preliminaries of concepts and definitions. Section 3 examines the ‘intrinsic’ case for inter-group equality, with an emphasis on the graphical device called the group poverty profile. Section 4 analyses the ‘instrumental’ case for inter-group equality, by exploring the possibility of a group-affiliation externality determining deprivation status at the individual level. Section 5 concludes.

2 Concepts and definitions

Let N be the set of positive integers, R the real line, and R_{++} the positive real line. For every $n \in N$, let \mathbf{X}_n be the set of non-negative n -vectors $\mathbf{x} = (x_1, \dots, x_i, \dots, x_n)$, where x_i is the income of individual i in a community of n individuals. Define the set $\mathbf{X} \equiv \cup_{n \in N} \mathbf{X}_n$. Let the poverty line be denoted by z , where z is a *positive and finite* income level such that all persons with income less than z are identified as being poor. For all $\mathbf{x} \in \mathbf{X}$, $N(\mathbf{x})$ will stand for the set of all individuals whose incomes are represented in the income distribution \mathbf{x} , and $Q(\mathbf{x})$ will stand for the set of poor individuals. Next—and drawing on Jayaraj and Subramanian (1999) and Subramanian (2004b)—for every $n \in N$, let \mathbf{G}_n be the set of all possible partitions of the set $N = \{1, \dots, n\}$, and define the set $\mathbf{G} \equiv \cup_{n \in N} \mathbf{G}_n$. Every $\mathbf{g} \in \mathbf{G}$ is some partition of the population, induced by some appropriate *grouping*, for example on the basis of race, caste, gender, etc.; and the elements of \mathbf{g} —denoted by the running index j —will be taken to be *subgroups* of the population. Clearly, for every $\mathbf{g} \in \mathbf{G}$ and $n \in N$, it will be the case that $1 \leq \#\mathbf{g} \leq n$. Two polar cases of grouping are the *atomistic grouping* \mathbf{g}^a which induces the finest partition $\{\{1\}, \dots, \{n\}\}$ of $\{1, \dots, n\}$, and the *universal grouping* \mathbf{g}^u which induces the grossest partition $\{\{1, \dots, n\}\}$ of $\{1, \dots, n\}$. For all $(\mathbf{x}, \mathbf{g}) \in \mathbf{X} \times \mathbf{G}$, the pair (\mathbf{x}, \mathbf{g}) will be said to be a *compatible* pair if and only if \mathbf{g} partitions a population whose size is the same as the dimensionality of \mathbf{x} . Given any compatible pair $(\mathbf{x}, \mathbf{g}) \in \mathbf{X} \times \mathbf{G}$, \mathbf{x}_j will stand for subgroup j 's income vector, and $\mu_j \equiv \mu(\mathbf{x}_j)$ will stand for the mean income of subgroup j , for every $j \in \mathbf{g}$. Further, for all compatible $(\mathbf{x}, \mathbf{g}) \in \mathbf{X} \times \mathbf{G}$, \mathbf{x}^i will stand for the income vector of the subgroup of \mathbf{g} to which person i belongs, and $\mu^i \equiv \mu(\mathbf{x}^i)$ will denote the average income of the subgroup to which i belongs, for every $i \in N(\mathbf{x})$. We now define a poverty index formally.

A *poverty index* is a mapping $P: R_{++} \times \mathbf{X} \times \mathbf{G} \rightarrow R$ such that, for every $z \in R_{++}$, and every compatible pair $(\mathbf{x}, \mathbf{g}) \in \mathbf{X} \times \mathbf{G}$, $P(z, \mathbf{x}, \mathbf{g})$ specifies a real number which is intended to represent the extent of poverty associated with the regime $(z, \mathbf{x}, \mathbf{g})$.

Certain standard axioms invoked in the measurement of poverty are now swiftly and informally reviewed. *Focus* (Axiom F) requires the poverty index to be invariant with respect to increases in non-poor incomes; *continuity* (Axiom C) requires the poverty index to be continuous on \mathbf{X}_n for every $n \in N$; *normalization* (Axiom N) requires the poverty index to attain a lower bound of zero when there is no poor person in the community; *symmetry* (Axiom S) requires the poverty index to be invariant to any interpersonal permutation of incomes; *monotonicity* (Axiom M) [respectively, *weak monotonicity* (Axiom WM)] requires the poverty measure to rise [respectively, not decline] with a decline in the income of any poor person; *respect for income dominance* (Axiom D; see Amiel and Cowell 1994, and Subramanian 2004b) requires that, *ceteris paribus*, poverty associated with the vector \mathbf{x} be lower than that associated with the vector \mathbf{y} whenever \mathbf{x} vector-dominates \mathbf{y} ; *transfer* (Axiom T) requires that the poverty index register a decline in value whenever, other things equal, there is a rank-preserving transfer of income from a poor person to a poorer person; *decomposability* (Axiom D) requires the poverty index to be expressible as a population-share weighted average of sub-group poverty levels; and *subgroup sensitivity* (Axiom SS; see Jayaraj and Subramanian 1999, and Subramanian 2004b) requires that, other things equal, poverty should decline whenever a pure redistribution of incomes between two groups causes the relatively disadvantaged group to become less poor and the relatively advantaged group to become poorer while maintaining the relative poverty rankings of the two groups.

A poverty index P is said to be *degenerate with respect to grouping* (or just *degenerate*) if, for every $z \in R_{++}$, every $\mathbf{x} \in \mathbf{X}$ and all distinct $\mathbf{g}, \mathbf{g}' \in G$ such that (\mathbf{x}, \mathbf{g}) and $(\mathbf{x}, \mathbf{g}')$ are compatible pairs, it is the case that $P(z, \mathbf{x}, \mathbf{g}) = P(z, \mathbf{x}, \mathbf{g}')$, that is, the value of the poverty index is invariant with respect to the grouping employed. For future reference, we shall denote by Π the set of poverty indices which are symmetric, weakly monotonic, decomposable, normalized to lie in the interval $[0, 1]$, and degenerate.

3 Poverty aggregation when inter-group equality is intrinsically valued

It can be shown that a poverty index which is required to simultaneously satisfy the sets of requirements constituted by {symmetry, monotonicity, subgroup sensitivity}, or {respect for income dominance, transfer, subgroup sensitivity}, could run into existence problems, on which see Subramanian 2004b: Propositions 4.1 and 4.2. In assessing these impossibility results, the author points to the problematic outcome of insisting on the universal validity of certain axioms irrespective of the context in which they are invoked. A reasonable way out might be to restrict the applicability of a given axiom to a domain that is relatively non-controversial. Thus, in the case of the symmetry axiom, one may require its brief to extend only to the extent that *within* any subgroup, swapping incomes across members of the subgroup should not alter the level of subgroup poverty, and that in a *between-group* context, the level of aggregate poverty should not vary with the way in which the subgroups are labeled. Similarly, with the transfer axiom, its sway could be limited to interpersonal income redistributions *within* a subgroup, while in a between-group context one could require that for a given level of poverty averaged across the subgroups, aggregate poverty should increase with an increase in inequality in the inter-group distribution of poverty. Poverty indices which satisfy these restricted versions of symmetry and transfer also possess the virtue of a certain sort of ‘flexibility’, in terms of which, for example, an inter-personal transfer across members of different subgroups which reduces inter-group disparity in the distribution of poverty may or may not reduce aggregate poverty, depending on how regressive the transfer is: a more rather than less regressive transfer could be partial to the transfer axiom at the expense of the subgroup sensitivity axiom, and the other way around with a less rather than more regressive transfer. Examples of such poverty indices are available in a poverty-related version of the Anand-Sen ‘gender-adjusted human development index’, and in the measure advanced by Jayraj and Subramanian (1999), and discussed in Subramanian and Majumdar (2002). These poverty measures are presented below.

First, let (\mathbf{x}, \mathbf{g}) be a compatible pair belonging to $\mathbf{X} \times G$, and let z be any poverty line. Let π be any poverty index belonging to the set Π of measures which are symmetric, weakly monotonic, decomposable, normalized to lie in the interval $[0, 1]$, and degenerate. We shall assume that \mathbf{g} partitions the population into K exclusive and exhaustive subgroups. $\boldsymbol{\pi}$ will stand for the *non-increasingly ordered* vector of subgroup poverty levels $(\pi_1, \dots, \pi_j, \dots, \pi_K)$, viz. $\boldsymbol{\pi}$ is ordered such that $\pi_j \geq \pi_{j+1}$ for all $j = 1, \dots, K - 1$. We let t_j stand for the population share of the j th poorest group, and $T_j \equiv \sum_{k=1}^j t_k$ for the cumulative proportion of the population belonging to groups whose poverty levels are greater than or equal to that of the j th group. For future reference, we also define, for every $j = 1, \dots, K$, the quantity $S_j \equiv \sum_{k=K-j+1}^K t_k$. Recall that since π is a decomposable index, it can be written as a population-share weighted sum of subgroup poverty levels: $\pi = \sum_{j=1}^K t_j \pi_j$. If all groups are of the same size, then it is easy to see that all subgroup poverty levels (the π_j) are accorded the same weight ($1/K$) by the poverty index π . To

secure an egalitarian tilt, one could think of a system of weights in which a higher weight is accorded to a subgroup with greater poverty. This is the basic idea underlying the construction of ‘group-inequality sensitive’ indices of aggregate poverty. A specific version of the Anand-Sen index (suitably adapted to the present context) is the index π^0 , and the poverty measure advanced in Subramanian and Majumdar (2002) is the index π^{00} , which—given z , \mathbf{x} and \mathbf{g} —can be written as follows:

$$(1) \pi^0(z, \mathbf{x}, \mathbf{g}) = [\sum_{j=1}^K t_j \pi_j^2]^{1/2}; \text{ and}$$

$$(2) \pi^{00}(z, \mathbf{x}, \mathbf{g}) = [1/(K-1)] \sum_{j=1}^K [(K-1-j)t_j + T_j] \pi_j.$$

If $C^* \equiv [(1/\pi^2) \sum_{j=1}^K t_j \pi_j^2] - 1$ is the squared coefficient of variation in the inter-group distribution of the π_j , and $G^* \equiv [1/\pi(K+1)] \sum_{j=1}^K [(K+1-j)t_j + T_j] \pi_j - 1$ is the Gini coefficient of inequality for this distribution, then (1) and (2) can be written, respectively, as

$$(1') \pi^0 = \pi(1+C^*)^{1/2}; \text{ and}$$

$$(2') \pi^{00} = \pi[1 + \{(K+1)/(K-1)\}G^*].$$

Note also that

$$(2'') \text{Lim}_{K \rightarrow \infty} \pi^{00} = \pi(1 + G^*).$$

The indices π^0 and π^{00} are essentially the average level of poverty (π) enhanced by a factor incorporating the extent of inequality in the distribution of the group-specific levels of poverty: in this sense, their construction is a reflection of an intrinsic preference for inter-group equality. As we have already seen, concerns for group identity and inter-group (as distinct from inter-personal) equality could sit uneasily with the requirements of canonical axioms like symmetry and transfer. These issues have been discussed in some detail in Subramanian (2004b), and will be reviewed, in the context of ‘externality-adjusted’ poverty indices, in the following section. For this reason, I will not here dwell on these themes. Rather, an alternative way of deriving (a close relative of) the index π^{00} will be discussed. This entails the use of a graphical device called the *group poverty profile*, in presenting which I shall very closely follow Shorrocks’ (1995, 1999) construction of the *poverty gap profile*.

Given the ordered vector $\boldsymbol{\pi}$ of subgroup poverty levels ($\pi_1, \dots, \pi_j, \dots, \pi_K$), the group poverty profile (GPP) is obtained as a plot of the points (D_0, T_0) , $(D_1, T_1), \dots, (D_j, T_j), \dots, (D_K, T_K)$, where $D_0 = T_0 \equiv 0$ and, for every $j = 1, \dots, K$:

$$(3) D_j(\boldsymbol{\pi}; T_j) = \sum_{k=1}^j t_k \pi_k.$$

The GPP is thus constructed by first arranging the subgroups in the order of poorest to least poor in terms of the values of the π_j ; the population-share weighted poverty levels are then cumulated across the subgroups and plotted against the cumulative population shares of the subgroups; and the graph obtained by connecting the plotted points by straight lines yields a ‘piece-wise linear’ version of the GPP. Since the maximum value the π_j can take is unity, it is clear that the ‘worst-case picture’ is obtained when $\pi_j = 1$ for all j , in which situation the GPP will be the diagonal of the unit square, and can be called the ‘line of maximal poverty’ (in all of which we are, with appropriate adaptation to the context, closely following Shorrocks, 1996). In general, it can be seen that the

GPP can be drawn as a non-decreasing, concave curve which lies beneath the diagonal of the unit square. The highest point on the GPP is obviously its final point (D_K, T_K) whose height is $\sum_{k=1}^K t_k \pi_k$, or π , recalling that π is a decomposable index. For illustrative purposes, a typical GPP is drawn in Figure 1, where K is taken, for specificity, to be 4.

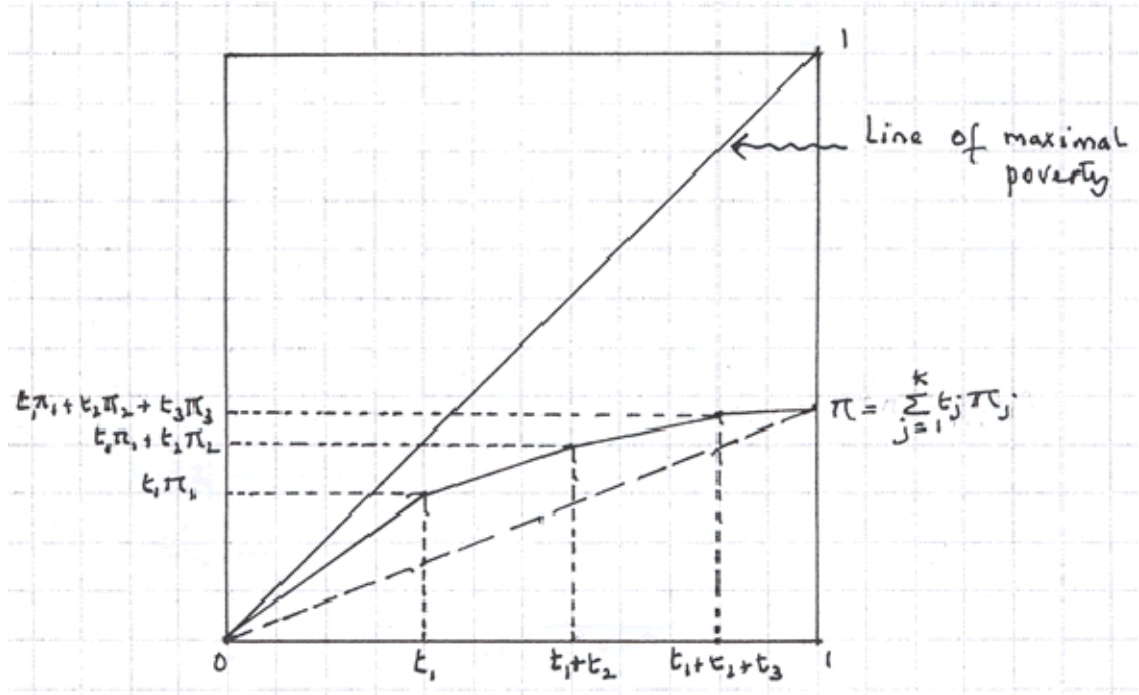


Figure 1 A group poverty profile

Notice from Figure 1 that if there were no inequality in the inter-group distribution of poverty levels, that is, if it were the case that $\pi_j = \pi$ for all j , then the GPP would be the straight broken line connecting the points 0 and π in the figure. The actual GPP lies above the broken line, and it is natural to attribute the space enclosed by the two curves to the fact of an unequal distribution of subgroup poverty levels. Indeed, an inverted image of the broken line and the GPP suggests something like a Lorenz curve drawn beneath the line of equality, and it is to a construction, precisely, of the *group poverty Lorenz profile* (GPLP) that we now turn. The GPLP is obtained from plotting the points $(L_0, S_0), (L_1, S_1), \dots, (L_j, S_j), \dots, (L_K, S_K)$, where $L_0 = S_0 \equiv 0$ and, for every $j = 1, \dots, K$:

$$(4) L_j(\pi; S_j) = (1/\pi) \sum_{k=K-j+1}^K t_k \pi_k.$$

The graph is drawn by first ranking the subgroups in *non-decreasing order* of their poverty levels, plotting the cumulative subgroup shares in total poverty against their cumulative population shares, and connecting the plotted points by straight lines to yield a ‘piece-wise linear’ GPLP. Figure 2 features a typical GPLP drawn within the unit square, for the special case in which $K = 4$.




Figure 2 A group poverty Lorenz profile

Continuing to follow the lead afforded in Shorrocks (1996), one can see that the group poverty profile and the group poverty Lorenz profile are linked by the following relationship:

$$(5) D_j(\boldsymbol{\pi}; T_j) = \pi[1 - L_{K-j}(\boldsymbol{\pi}; S_{K-j})], j = 0, 1, \dots, K.$$

From (3), (4) and (5) one can see that the lower the GPP or the higher the GPLP is, the further away will the GPP be from the line of maximal poverty, and so the ‘better’ may we judge the poverty situation to be. This paves the way for comparing alternative group poverty distributions in terms of a *poverty dominance* relationship based on the GPP. Specifically, for any given poverty line and grouping which partitions the population into K subgroups, we let $\boldsymbol{\pi}$ (as before) stand for the non-increasingly ordered vector of subgroup poverty levels (π_1, \dots, π_K) and $\mathbf{t} = (t_1, \dots, t_K)$ for the corresponding vector of subgroup population shares. Given $\boldsymbol{\pi}$ and \mathbf{t} , a *poverty situation* \mathbf{s} is defined as a K -tuple of pairs of subgroup poverty level and subgroup population share: $\mathbf{s} = ((\pi_1, t_1), \dots, (\pi_K, t_K))$. For any two poverty situations $\mathbf{s}^1 = ((\pi^1_1, t^1_1), \dots, (\pi^1_K, t^1_K))$ and $\mathbf{s}^2 = ((\pi^2_1, t^2_1), \dots, (\pi^2_K, t^2_K))$, we shall say that \mathbf{s}^1 *poverty-wisely dominates* \mathbf{s}^2 , written as $\mathbf{s}^1 >_p \mathbf{s}^2$, whenever it is the case that the GPP for \mathbf{s}^1 lies somewhere below and nowhere above the GPP for \mathbf{s}^2 . (Of course, when subgroup sizes are different in the two poverty situations under comparison, the contributions of population size variations and subgroup poverty variations to the difference in aggregate poverty between the two situations would need to be unscrambled.) In general, the GPP (much like the Lorenz curve employed in standard distributional analysis) is a useful visual aid to group-related poverty analysis, and leads naturally to the construction of a strict partial ordering of poverty, namely the binary

relation \succ_p . It could find particularly fruitful application in the over-time assessment of, say, the spatial distribution of poverty.

Finally, the GPP also paves the way for deriving a measure of aggregate poverty, in much the same way as the Lorenz curve paves the way for deriving the Gini coefficient of inequality. Specifically—and still closely tracking Shorrocks (1995, 1996)—it appears to be natural to obtain a normalized measure of poverty—call it π^* —in terms of the area beneath the GPP expressed as a proportion of the area beneath the line of maximal poverty (recall Figure 1). The area beneath the line of maximal poverty (call this Area A) is just one-half, while the area beneath the GPP (call this Area B), obtained as a sum of the areas of a number of trapezoids, can—with some manipulation—be seen to be given by: $\text{Area B} = \sum_{j=1}^K t_j T_j \pi_j - \sum_{j=1}^K t_j \pi_j^2 / 2$. Then, the poverty index π^* (= Area B/Area A) will be given by

$$(5) \pi^* = 2\sum_{j=1}^K t_j T_j \pi_j - \sum_{j=1}^K t_j \pi_j^2.$$

Turning next to Figure 2, and noting that the Gini coefficient of inequality G^* in the distribution of subgroup poverty levels is just the area beneath the GPLP divided by the area beneath the diagonal of the unit square, it is easy to verify that

$$(6) G^* = (2/\pi)[\sum_{j=1}^K t_j T_j \pi_j - \sum_{j=1}^K t_j \pi_j^2 / 2] - 1.$$

Making the appropriate substitution from (6) into (5) yields:

$$(7) \pi^* = \pi(1 + G^*).$$

Comparing (7) with (2'') reveals that the index π^* is just the asymptotic version of the index π^{00} .

Suppose the grouping of the population we resorted to were the ‘atomistic grouping’, in terms of which each person is regarded as constituting a group by herself/himself, so that a typical group— $\{i\}$ —is the one constituted by person i , for every $i \in N$. Suppose further that x_i is the income of the i th poorest person and that, for every i , $\pi_{\{i\}}$ is given by: $\pi_{\{i\}} = \max\{0, 1 - x_i/z\}$, that is, a person’s deprivation is measured by the proportionate shortfall of her income from the poverty line if she happens to be poor, and is taken to be zero otherwise. Under these circumstances, the GPP will be precisely what Shorrocks (1995, 1996) calls the *poverty gap profile*, and the index of aggregate poverty corresponding to π^* in (7) will be an index Δ which is given by $\Delta = HI(1 + G_p)$, where H is the familiar headcount ratio (or proportion of the population in poverty), I is the income-gap ratio (or proportionate shortfall of the average income of the poor from the poverty line), and G_p is the Gini coefficient of inequality in the distribution of the poverty gap values. Δ , as it happens, is precisely the poverty index derived by Shorrocks (1996), which he shows to be closely related to Sen’s (1976) poverty index, and even more proximately so to Thon’s 1979 index, of which Δ is the replication-invariant ‘asymptotic’ version. Of the index Δ , Shorrocks says:

The index obtained by taking the area below the poverty gap profile relative to that below the ‘line of maximum poverty’ has a particularly appealing interpretation. Although not decomposable across population subgroups, it is an ideal poverty index in all other respects. Shorrocks (1996: 251)

That there is a way of reckoning subgroup poverty in the measurement of aggregate poverty which is, substantially, a generalization of a particularly instructive approach to poverty measurement as advanced by Shorrocks is, we believe, both interesting and useful. We turn now to an instrumental justification for taking seriously the group-wise distribution of poverty in the aggregation exercise.

4 Poverty aggregation when inter-group equality is instrumentally valued

4.1 Motivation

This section explores a simple idea on the connection between individual and group poverty, and examines some straightforward implications of the idea for the measurement of aggregate poverty. The idea in question is this: in principle, how poor a person is, is dependent not only on how deprived s/he is, considered as an atomistic unit, but also on the general level of prosperity of the group to which s/he is affiliated. A good deal of the literature on poverty measurement rests on the view of man as economic man, i.e. as an entity wrenched loose from his group moorings. This section considers some consequences of taking a view of man as social animal, i.e. as an entity whose fortunes are tied, in smaller or larger measure, to the fact of his particular group affiliation. Within this setting, certain conventional approaches to, and outcomes of, poverty measurement can be seen to emerge as special cases of a more general framework that allows for intra-group externality in the determination of individual deprivation.

The issue of how outcomes at the individual level are influenced by the fact of group affiliation is well brought out, at a general level, in Loury (2000: 233, 243):

Economic analysis begins with a depersonalized agent who acts more or less independently to make the best of the opportunities at hand ... [But] [e]ach individual is socially situated, and one's location within the network of social affiliation substantially affects one's access to various resources ... There is one view of society in which we are atomistic individuals, pursuing our paths to the best of our abilities ... But this is a false, or at least incomplete, view of how society works. The fact is that we are all embedded in a complex web of associations, networks, and contacts. We live in families, we belong to communities, and we are members of collectivities of one kind or another. We are influenced by these associations from the day we are born. Our development—what and who we are and become—is nourished by these associations.

At a more specific level—that of caste in the Indian context—B. R. Ambedkar (cited in Guhan 2000: 206) makes a pointed observation about intra-group externality when he says:

Take, for instance, Chamars, you look upon this community with hatred, but if there are some lawyers, doctors and educated persons among them, you cannot put your hand upon them and you will not do that, although everyone of them is not so highly educated. You will say he is a Bhangi

but suppose there are educated persons among them, you will respect them.

The thrust of the preceding considerations, when applied to the context of poverty, resolves itself into a straightforward principle which constitutes the underlying motivation for the present section, and can be stated in terms of the following axiom:

Axiom of Group-Mediated Deprivation (Axiom GMD). *Ceteris paribus*, of two poor persons having the same income, the person who belongs to the group with a lower mean income is more deprived.

The idea underlying Axiom GMD is that a person's ability to transform income into what Sen (1985) calls functionings is linked to the fortunes of the group to which the person is affiliated. This type of intra-group externality could occur in a number of ways. It is often the case that the extent to which a poor individual is a beneficiary of 'informal social security', or of credit facilities, or of supplementary nutrition, or of assistance in the event of sickness or disability, or of scholarships in educational institutions, is an increasing function of the average level of income sufficiency of the group to which the person is affiliated. If a poor person's access to resources varies directly with the respect and consideration with which s/he is treated by society at large, then such respect and consideration, and therefore personal access to resources, might be expected to be positively correlated with the mean prosperity of the group to which the person belongs. Moreover, a person's own sense of advantage is often mediated by the economic status of the group of which s/he is a member: even in the absence of differentials in income at the inter-personal level, one can expect inter-personal differentials in the level of achieved functionings arising from group-mediated individual assertiveness, or diffidence, as the case may be.

In much of what follows, *illustrative* operational content is given to Axiom GMD. The idea is not to claim any particular social realism for the specific formulations pressed into service; rather, it is to draw out certain elementary conceptual implications of the axiom for poverty comparisons, poverty axiomatics, and related settings of measurement.

4.2 An 'externality-adjusted' poverty measure

Consider any $(z, \mathbf{x}, \mathbf{g}) \in R_{++} \times \mathbf{X} \times \mathbf{G}$, where (\mathbf{x}, \mathbf{g}) is a compatible pair. For any person i , let d_i represent i 's *deprivation function*. It is reasonable to suggest that, for all $i \in \{1, \dots, n\}$:

$$d_i = d_i(x_i, \mu^i)$$

that is, i 's deprivation level depends on both his own income and the average income of the group to which he belongs. Further, we could imagine that d_i is declining in each of x_i and μ^i : a person's deprivation declines as her income increases, and likewise declines as the average income of the group to which she is affiliated increases, which captures the 'group externality effect' postulated by the axiom of group-mediated deprivation. Additionally, and with an eye to normalization, we may suppose that $d_i(0,0) = 1$, and $\lim_{\mu^i \rightarrow \infty} d_i(x_i, \mu^i) = 0$: the poverty of a completely poor person belonging to a completely poor group is unity, while a person who is poor by virtue of his income falling short of the poverty line will be taken to approximate to the status of non-poorness as the

average income of the group to which he belongs becomes indefinitely large. Finally, a person with income equalling or exceeding the poverty line will be taken to be non-poor, namely $d_i(x_i, \mu^i) = 0$ for $x_i \geq z$. A functional form for d_i which satisfies all of the requirements stated above is given by

$$d_i(x_i, \mu^i) = (z - x_i)/(z + \mu^i).$$

The particular specialization resorted to above is not sought to be uniquely characterized: it is employed only as a convenient illustrative device. I shall now define an ‘externality-adjusted’ poverty index, P^* , as a simple average of the individual-specific deprivation functions, viz. for all $z \in R_{++}$, and every compatible pair $(\mathbf{x}, \mathbf{g}) \in \mathbf{X} \times \mathbf{G}$:

$$(8) P^*(z, \mathbf{x}, \mathbf{g}) [= (1/n) \sum_{i \in Q} d_i(x_i, \mu^i)] = (1/n) \sum_{i \in Q} [(z - x_i)/(z + \mu^i)]$$

where Q is the set of poor individuals.

(8) can also be written as

$$(9) P^*(z, \mathbf{x}, \mathbf{g}) = (1/n) \sum_{j \in \mathbf{g}} [\sum_{i \in Q^j} \{(z - x_i)/(z + \mu_j)\}]$$

where Q^j is the set of poor persons belonging to subgroup j . Notice now that if P^*_j is a shorthand for $P^*(z, \mathbf{x}_j, \mathbf{g}^j)$, namely if P^*_j is the poverty level of the j th subgroup when the partitioning of the subgroup’s population is the universal one, then

$$(10) P^*_j = (1/n_j) \sum_{i \in Q^j} \{(z - x_i)/(z + \mu_j)\}$$

where n_j is the dimensionality of \mathbf{x}_j ; whence—in view of (9):

$$(11) P^*(z, \mathbf{x}, \mathbf{g}) = (1/n) \sum_{j \in \mathbf{g}} n_j P^*_j = \sum_{j \in \mathbf{g}} t_j P^*_j$$

where $t_j \equiv n_j/n$ is the population share of the j th subgroup.

The two polar cases of grouping we have considered earlier are of special interest. When the grouping is atomistic, that is, each person is considered to constitute a group by himself or herself, then clearly $\mu^i = x_i$ for all i , whence, in view of (8), we have:

$$(12) P^*(z, \mathbf{x}, \mathbf{g}^a) = (1/n) \sum_{i \in Q} [(z - x_i)/(z + x_i)].$$

(12) is the sort of ‘individualistic’ poverty index most widely employed in the poverty measurement literature: it turns out to be a special case of the more general formulation (8). A matter of some interest is that this poverty index can be independently derived as a certain kind of normalized distance function, in the following sense. Given a non-decreasingly ordered income vector $\mathbf{x} = (x_1, \dots, x_q, x_{q+1}, \dots, x_n)$, where $x_q < z$ and $x_{q+1} \geq z$, define the *censored* version \mathbf{x}^c of \mathbf{x} as the vector obtained by replacing all the non-poor incomes in \mathbf{x} by the poverty line income z (see Takayama 1979), so that $\mathbf{x}^c = (x_1, \dots, x_q, z, \dots, z)$. Let \mathbf{z} be the n -vector (z, \dots, z) and $\mathbf{0}$ the n -vector $(0, \dots, 0)$. Then (see Subramanian 2004a), \mathbf{z} can be interpreted as the income distribution with the smallest mean which is compatible with a complete absence of poverty, and $\mathbf{0}$ as the income distribution representing total poverty. Let $\delta(\mathbf{z}, \mathbf{x}^c)$ represent the ‘vector distance’ between \mathbf{z} and \mathbf{x}^c (the shortfall of the ‘actual’ situation from the ‘no-poverty’ situation), and $\delta(\mathbf{z}, \mathbf{0})$ represent the vector distance between \mathbf{z} and $\mathbf{0}$ (the shortfall of the ‘complete

poverty' situation from the 'no-poverty' situation). Then, in some intuitively straightforward sense we could take the ratio of the two distances, call it $r \equiv \delta(\mathbf{z}, \mathbf{x}^c) / \delta(\mathbf{z}, \mathbf{0})$, to be a normalized measure of poverty. It remains to specify the form of the distance function δ . One candidate is the *Camberra distance function* δ^C (see Wilson and Martinez 1997) which calculates the distance between any two n -vectors \mathbf{a} and \mathbf{b} as: $\delta^C(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^n [(a_i - b_i) / (a_i + b_i)]$. This leads to:

$$(13) r^C \equiv \delta^C(\mathbf{z}, \mathbf{x}^c) / \delta(\mathbf{z}, \mathbf{0}) = [\sum_{i \in Q} \{(z - x_i) / (z + x_i)\}] / [\sum_{i=1}^n (z / z)] \\ = (1/n) \sum_{i \in Q} [(z - x_i) / (z + x_i)].$$

From (12) and (13) we note that the measure $P^*(z, \mathbf{x}, \mathbf{g}^a)$ is just the normalized Camberra distance ratio r^C —which confers a simple and handy interpretation on the poverty index.

At the other polar extreme, when $\mathbf{g} = \mathbf{g}^u$, it can be seen from (8) that

$$(14) P^*(z, \mathbf{x}, \mathbf{g}^u) = (1/n) \sum_{i \in Q} [(z - x_i) / (z + \mu)]$$

where μ is the mean of the distribution \mathbf{x} . This corresponds to the case where there is only one group—that constituted by the grand coalition of individuals: the community, here, is regarded as being entirely homogeneous. (12) and (14), it may be reiterated, are special cases of (8).

4.3 Some implications of 'grouping' for poverty rankings and poverty axiomatics

4.3.1 Poverty rankings

The particular manner in which we choose to partition the population can have non-trivial implications for poverty comparisons. This proposition can be illustrated by means of a simple numerical example. Consider two 5-vectors of income, given, respectively, by $\mathbf{x} = (10, 20, 30, 50, 70)$ and $\mathbf{y} = (10, 20, 30, 45, 76)$. Let the poverty line z be 40. Consider also three alternative groupings of the population: the atomistic grouping \mathbf{g}^a , the universal grouping \mathbf{g}^u , and a grouping \mathbf{g}^r by religion which, let us say, divides the population into two religious groups 1 and 2 respectively, and precipitates the sub-group income vectors $\mathbf{x}_1 = (10, 50)$, $\mathbf{x}_2 = (20, 30, 70)$, $\mathbf{y}_1 = (10, 45)$, and $\mathbf{y}_2 = (20, 30, 76)$. Routine computation, employing equations (12), (14) and (8) respectively, will yield the following results:

$$P^*(z, \mathbf{x}, \mathbf{g}^a) = P^*(z, \mathbf{y}, \mathbf{g}^a) = .2152;$$

$$P^*(z, \mathbf{x}, \mathbf{g}^c) = .1579 > P^*(z, \mathbf{y}, \mathbf{g}^c) = .1575; \text{ and}$$

$$P^*(z, \mathbf{x}, \mathbf{g}^r) = .1607 < P^*(z, \mathbf{y}, \mathbf{g}^r) = .1621.$$

For an atomistic grouping of the population, the extent of poverty, as measured by the index P^* , is the same for both vectors \mathbf{x} and \mathbf{y} ; for the universal grouping, P^* suggests that poverty is greater in \mathbf{x} than in \mathbf{y} ; and for a grouping according to religion, P^* certifies that there is more poverty in \mathbf{y} than in \mathbf{x} . These pairwise rank-reversals indicate that it clearly cannot be a matter of indifference how we choose to partition the population. In particular, and almost in a spirit of absent-mindedness, the only partitioning which has generally been held to be relevant is the atomistic one—to the

point that it has conventionally not even been considered to be important to specify the particular partitioning invoked as an argument in the poverty function. But some concern for how sociology could refine one's understanding of economics would necessarily point to the requirement of engaging actively with the appositeness of the particular grouping one resorts to in effecting poverty comparisons. The fact of intra-group externality induces a heterogeneity between groups that compels the need for a context-based attentiveness to the social dimension of economic theorizing. That this issue could also have practical implications for the *targeting* of anti-poverty budgetary allocations is straightforward (see Jayaraj and Subramanian 1999, and Subramanian 2004b).

4.3.2 *Poverty axiomatics*

What are some of the properties of the poverty index P^* , in relation to the 'standard' axioms of poverty measurement reviewed in Section 2? It is fairly easy to see that P^* satisfies the monotonicity and normalization axioms; and writing the index in the form of equation (12) suggests that it is also decomposable. These may well be the limits of P^* 's 'successes'. The focus axiom is violated by P^* : the numerical example reviewed in the preceding subsection indicates that the vector y has been derived from the vector x by changes to the non-poor incomes which have left the numbers of individuals in poverty unchanged; yet for both the universal grouping and the grouping according to religion, the value of P^* is *not* the same for x and y . The symmetry axiom could also be a casualty. If two persons with distinct incomes and belonging to different groups were to swap their incomes, then the group-specific means would change, and—see equation (10)—the value of P^* could also change. It is not hard to perceive that P^* could also fall foul of the transfer axiom. A progressive transfer of income from a poor person belonging to a relatively badly-off group to a poorer person belonging to a relatively well-off group could change group-specific means in such a way as to actually cause the value of P^* to *rise*, in opposition to what the transfer axiom demands. These issues are discussed more elaborately in Jayaraj and Subramanian (1999) and Subramanian (2004b). The point to note is that these 'axiomatic failures' of an index such as P^* are not necessarily adverse reflections on P^* : they may simply be a reflection of the inappropriateness of certain canonical normative properties of poverty measures for contexts which demand that we take the notion of groups and social heterogeneity seriously. 'Compensatory discrimination', 'affirmative action', and similar principles of group-mediated justice would be impossible to defend if we insisted on swearing context-independent allegiance to axioms like focus, symmetry and transfer.

4.3.3 *An application to a non-income dimension: measuring literacy*

Basu and Foster (1998) have shown how the conventional headcount measure of literacy could be a misleading indicator of a population's 'effective' literacy status, simply because of the failure of this measure to reckon the intra-household externality accruing to illiterate members of a household from the literacy of their literate cohorts. The externality-related approach to measuring income-poverty outlined in this paper has an intimate implication for the Basu-Foster thesis, which is elaborated in what follows. Consider a population of n individuals, partitioned into K households. A literate person will be taken to have a literacy status of unity, and an illiterate person a literacy status of zero. Corresponding to the income-poverty line z , we shall specify a 'literacy-poverty line' of unity. Suppose m persons in the n -person population are literate. Then the

‘standard’ measure of literacy, R , would be given by the headcount ratio m/n ; or, since the context of discussion here is deprivation rather than achievement, the standard measure of *illiteracy* is given by $R^0 \equiv 1 - R$. This measure is invariant with respect to the precise manner in which the m literates are distributed across the K households. If one allowed for a beneficent intra-household externality conferred by the literate members of a household on its illiterate members, then one can see that ‘effective literacy’ would rise (or ‘effective illiteracy’ would decline) with a more equitable inter-household distribution of literates. Hence Basu and Foster’s *instrumental* justification, based on an argument of efficiency-promoting externality, for planned literacy programmes that aim at equity in the inter-household distribution of literates. Precisely the same outcome is secured if illiteracy were to be measured analogously to the way in which income-poverty is measured by the index P^* . To see this, it is useful to refer to equations (9), (10) and (11).

In the context of illiteracy measurement, we replace z by unity; and ‘ x_i ’ for every illiterate person i is zero. For every household j belonging to the set $\{1, \dots, K\}$ of households, the ‘average’ literacy level is represented by the household’s literacy rate R_j (the proportion of literates in the j th household). For every j , Q^j is the set of illiterates in household j . It can be noted now from (10) that, if n_j is the size of household j , and if q_j is the cardinality of Q^j , then the ‘externality-adjusted’ or ‘effective’ illiteracy rate of household j is given by

$$\begin{aligned} R_j^{0*} &= (1/n_j) \sum_{i \in Q^j} \{(1 - 0)/(1 + R_j)\} = (q_j/n_j)/(1 + R_j) \\ &= (1 - R_j)/(1 + R_j) \equiv R_j^0/(2 - R_j^0), \end{aligned}$$

where $R_j^0 \equiv (1 - R_j)$ is the proportion of illiterates in household j .

If ϕ_j is the population share of household j , then, in view of (11), the ‘effective’ society-wide illiteracy rate can be written as

$$(15) \quad R^{0*} [= \sum_{j=1}^K \phi_j R_j^{0*}] = \sum_{j=1}^K \phi_j \{R_j^0/(2 - R_j^0)\}.$$

Suppose now that the ‘crude’ headcount measure of illiteracy for some community is R^0 . We consider two situations: in situation 1 the inter-household distribution of (il)literacy is an equal one, viz. $R_j^0 = R^0$ for all $j = 1, \dots, K$. In situation 2, we have a completely polarized outcome: all illiterates belong to households without any literates, and all literates belong to households without any illiterates. Let the ‘effective’ illiteracy rates corresponding to these two situations be denoted, respectively, by $R^{0*}(1)$ and $R^{0*}(2)$. In view of (15), it is easily verified that

$$(16) \quad R^{0*}(1) = R^0/(2 - R^0); \text{ and}$$

$$(16') \quad R^{0*}(2) = R^0.$$

Indeed, between them, (16) and (16’) describe the lower and upper limits of the values which R^{0*} can attain for any given R^0 : $R^{0*} \in [R^0/(2 - R^0), R^0]$. Thus if, for example, 50 per cent of a population is illiterate, then under situation 1, $R^{0*} = 0.33$, whereas under situation 2, $R^{0*} = 0.50$. The customarily-employed ‘crude’ illiteracy rate, at 0.5 under both situations 1 and 2, will be unable to distinguish between the two situations, unlike the ‘externality-adjusted’ illiteracy rate R^{0*} , which favours the equitable inter-

household distribution over the inequitable one—precisely the point made by Basu and Foster.

Briefly, the phenomenon of intra-group externality is a pervasive one, and is fruitfully applied in the measurement of deprivation in more than one context and in more than one dimension.

5 Concluding observations

Horizontal, or inter-group, inequalities in the distribution of poverty are a standard feature of many stratified societies. Intrinsic considerations of group fairness would dictate a concern for incorporating inequalities in the group-wise distribution of poverty directly into the aggregation exercise of poverty measurement. Extant approaches to such aggregation procedures have been reviewed in this paper, and in the process, a generalization of Shorrocks' illuminating approach to poverty measurement, involving the derivation of 'deprivation indices' from 'deprivation profiles', has been provided. Additionally, inter-group equality in the distribution of poverty could be justified on instrumental grounds, from the recognition that the goals of equity and efficiency could become congruent because of the presence of group-related externalities and their role in determining deprivation at the individual level. Indeed, externalities arising from group affiliation are an integral and obvious aspect of everyday social existence, yet with few exceptions they would appear to have been largely neglected in the economics of measuring deprivation. In this paper, an attempt has been made to incorporate the social embeddedness of individuals into an exercise in measuring income-poverty, and—along lines motivated by Basu and Foster (1998)—illiteracy.

The concerns of this paper have been primarily methodological, and accordingly, the points it addresses have been sought to be made illustratively rather than definitively. A principal reason for this—especially in the context of an instrumental justification for taking group inequalities seriously—is the difficulty of actually measuring external effects in any precise and unambiguous way, which could well be the reason for their general neglect in the literature. Allowing for an explicit reckoning of inter-group inequalities in a heterogeneous population, or for group-mediated deprivation, enables one to see that how one partitions the population has implications for deprivation comparisons, for the appeal of normative properties of deprivation indices that have been customarily held to be self-evidently desirable, for practical stratagems of budgetary allocation toward redress of deprivation, and for an instrumental defence of inter-group equality and 'reverse' discrimination. The standardly 'homogenising', or 'atomistic', formulations of measurement concerns have been shown to emerge as special cases of a more general framework of group-inclusive analysis. Yet, it is the 'special case' of mainstream, or 'individualistic', economic theorizing which has held sway as the 'general' or 'paradigmatic' model, against which more socially realistic approaches have had to be judged as departures from the norm. As has been stated earlier, it is perhaps understandable that measurement exercises are largely confined to that which is more, rather than less, tractably measured. But a re-appraisal is also necessitated when these severely practical considerations have become overwhelmingly successful in preventing conceptual and normative aspects of social reality from infecting economic theorizing.

References

- Amiel, Y. and F. Cowell (1994) 'Monotonicity, Dominance and the Pareto Principle', *Economics Letters* 45: 447-50.
- Anand, S. and A. Sen (1995) 'Gender Inequality in Human Development: Theories and Measurement', Occasional Paper 19, Human Development Report Office. New York: UNDP. Cited in UNDP (1995).
- Basu, K. and J. E. Foster (1998) 'On Measuring Literacy', *Economic Journal* 108: 1733-49.
- Guhan, S. (2000) *Selected Writings of S. Guhan: India's Development Experience* (S. Subramanian, ed.), Delhi: Oxford University Press.
- Jayaraj, D. and S. Subramanian (1999) 'Poverty and Discrimination: Measurement, and Evidence from Rural India', in B. Harriss-White and S. Subramanian (eds), *Illfare in India: Essays on India's Social Sector in Honour of S. Guhan*, New Delhi: Sage Publishers.
- Jayaraj, D. and S. Subramanian (2000) 'Group Affiliation, Externality, and the Measurement of Poverty', *S. Guhan Memorial Series Discussion Paper 1*, Madras Institute of Development Studies: Chennai.
- Loury, G. (2000) 'Social Exclusion and Ethnic Groups: The Challenge to Economics', *Discussion Papers 106 (May)*, Institute for Economic Development, Boston University.
- Majumdar, M. and S. Subramanian (2001) 'Capability Failure and Group Disparities: Some Evidence from India for the 1980s', *Journal of Development Studies* 37: 104-140.
- Sen, A. K. (1976) 'Poverty: An Ordinal Approach to Measurement', *Econometrica* 44: 219-31.
- Sen, A. K. (1985) *Commodities and Capabilities*, Amsterdam and New York: North Holland.
- Shorrocks, A. F. (1995) 'Revisiting the Sen Poverty Index', *Econometrica* 50: 3-17.
- Shorrocks, A. F. (1996) 'Deprivation Profiles and Deprivation Indices', in S. Jenkins, A. Kapteyn and B. Vaan Praag (eds) *The Distributions of Welfare and Household Production: International Perspectives*. London: Cambridge University Press.
- Subramanian, S. (2004a) 'A Re-scaled Version of the Foster-Greer-Thorbecke Poverty Indices Based on an Association with the Minkowski Distance Function', *Research Paper 2004/10*, UNU-WIDER: Helsinki.
- Subramanian, S. (2004b) 'Social Groups and Economic Poverty: A Problem in Measurement', *Research Paper 2004/59*, UNU-WIDER: Helsinki.
- Subramanian, S. and M. Majumdar (2002) 'On Measuring Deprivation Adjusted for Group Disparities', *Social Choice and Welfare* 19: 265-80.
- Takayama, N. (1979) 'Poverty, Income Inequality, and their Measures: Professor Sen's Axiomatic Approach Reconsidered', *Econometrica* 47 (3): 747-59.

Thon, D. (1979) 'On Measuring Poverty', *Review of Income and Wealth* 25: 429-39.

United Nations Development Programme (UNDP) (1995) *Human Development Report 1995*, New York: Oxford University Press.

Wilson, D. R. and T. R. Martinez (1997) 'Improved Heterogeneous Distance Functions', *Journal of Artificial Intelligence Research* 6: 1-34.