Research Paper No. 2005/62

Poverty Measurement and Theories of Beneficence

S. Subramanian*

November 2005

Abstract

This note points to certain similarities of orientation and outcome between Derek Parfit's quest for a theory of beneficence and Amartya Sen's quest for a suitable real-valued representation of poverty. It suggests that both projects, in a certain sense, have been instructive failures. Using Sen's own work, the note also suggests a logically natural way of dealing with some of the problems in poverty measurement reviewed in it—but only to reject this way out on other compelling grounds.

Keywords: quantity, quality, inequality, identification, aggregation

JEL classification: I32

---

*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 92-9190-746-4 (internet version)

Acknowledgements

I am grateful to Tony Addison for helpful comments, and to D. Jayaraj for numerous discussions of the subject of this paper. The usual caveat applies.

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

Camera-ready typescript prepared by Lorraine Telfer-Taivainen at UNU-WIDER

1 Introduction

Amartya Sen and Derek Parfit have made two major contributions to the history of ideas—in the form, respectively, of an approach to the measurement of poverty, and the quest for a theory of beneficence, or human wellbeing. The present note makes a simple point: that there are similarities between the two enterprises which bear remarking. The commonalities, it is worth underlining, must not be overdrawn; and there is, in particular, no profit to be had in seeking a one-to-one correspondence (which does not exist) between the two strands of thought. It is, nevertheless, instructive to discern a certain broad identity of concerns shared by the exercises of measuring poverty and outlining the requisites of a theory of beneficence. The latter has implications for the former which are rather obvious, once they have been identified. The principal excuse for the writing of this note is that, to the best of this author's knowledge, these implications have not so far been highlighted.

Parfit's concerns arose from a consideration of certain problems in population ethics, including the one posed by the question: how many people should there ever be? In his book *Reasons and Persons* (hereafter *RAP*), he says (Parfit 1984: 381): 'In a complete moral theory, we cannot avoid this awesome question'. And an answer, he concluded, would depend upon the prospects of discovering 'Theory X', which is an acceptable theory of beneficence, or wellbeing, as reflected in '…the level of happiness, or … the quality of life, or … the share per person of resources. We should assume that, in my examples, these three correlate, rising and falling together'. An 'acceptable' theory of beneficence, from a reading of Parfit, would appear to be one which pays due attention to the *quantity* of wellbeing, the *quality* of wellbeing, and—where there are interpersonal differences in attained levels of wellbeing—the *extent of inequality* in its distribution. Thus, on page 405 of *RAP*, he says, 'I believe that the best theory about beneficence must claim that quality and quantity both have value'; and in subsequent discussions, particularly Chapter 19 of *RAP*, we discern considerable engagement with the question of inequality, where the suggestion is that Theory X should be informed by the principles of both beneficence *and* equality. How should Theory X combine considerations of quantity, quality and equality? In general, in such a way as not to offend our moral intuition, and in particular in such a way as to deliver us from certain Embarrassing Conclusions (of which one, in Parfit's nomenclature, is Repugnant and the other Absurd). For the purposes of this note, we do not need to get into the details of these conclusions. Simple numerical examples, as we shall see, should suffice to expose certain underlying difficulties, as we go along. Nevertheless, it may be useful to have a self-contained account of the problem of combining quantity and quality meaningfully in a calculus of beneficence—accordingly, a brief treatment of Parfit's concerns in this regard is provided in an Appendix.

There are parallels between the search for Theory X and the quest for a solution to the aggregation problem in poverty measurement, that is, the search for a 'satisfactory' measure of poverty with a real-valued representation. The parallels are strikingly evident from Sen's treatment of the problem in his classic *Econometrica* paper 'Poverty: An Ordinal Approach to Measurement', which—discounting Watts' (1968) seminal contribution—is in many ways the first systematic engagement with the question, and has since spawned an immense literature on the subject. It is instructive to read Sen's own interpretation of his poverty index (Sen 1976a, 1997: 227-8):

The poverty index proposed here turns out to have quite an easy interpretation. The measure [$P_S = H\{I + (1\text{-}I)G\}$] is made up of the headcount ratio H multiplied by the income gap ratio I augmented by the Gini coefficient G of the distribution of income among the poor weighted by (1-I), that is weighted by the ratio of the mean income of the poor to the poverty line income level. One way of understanding its rationale is the following: I represents poverty as measured by the proportionate gap between the mean income of the poor and the poverty line income. It ignores distribution *among* the poor, and G provides this information. In addition to the poverty gap of the mean income of the poor reflected in I, there is the 'gap' arising from the unequal distribution of the mean income, which is reflected by the Gini coefficient G of that distribution multiplied by the mean income ratio. The income-gap measure thus augmented to take note of inequality among the poor, that is I + (1-I)G, is normalized per poor person, and does not take note of the number of people below the poverty line, which could be minute or large. Multiplying [I + (1-I)G] by the headcount ratio H now produces the composite measure [P_S].

Poverty is about 'illbeing' (maleficence?), and is measured in the space of incomes, but setting aside these superficial contrasts, a glance at the index P_S suggests a strong congruence of concerns between Sen and Parfit. HI may be taken to represent a measure of the *quantity* of deprivation, I a measure of its *quality*, and G a measure of *inequality* in its interpersonal distribution. The ingredients are the same as those of Theory X; and the various extensions, refinements, and modifications of the Sen index of poverty are essentially variations on this theme.

If there are commonalities in the motivations underlying the Sen and Parfit programmes, what can one say of their respective outcomes? Similarities in motivation are, at best, an interesting *curiosum*; but similarities in outcome have rather more substantive implications. Parfit's pursuit of Theory X ended, by his own admission, in failure: as he says (Parfit 1984: 443), '[t]hough I failed to find such a theory, I believe that, if they tried, others could succeed'. Is the outcome of the poverty aggregation programme, such as it has evolved, a success? I shall argue, with the help of very elementary examples, against this notion; and in particular that the commonalities in orientation between the Sen and the Parfit programmes have carried over to their respective outcomes. At the same time, I shall also suggest that Sen's own work on a distributionally sensitive measure of 'real national income' may plausibly be exploited in the cause of measuring poverty in a more satisfactory way than has hitherto been obtained. Unfortunately, a slightly deeper scrutiny of this mode of resolution suggests—thanks again, as it happens, to Sen's own insights—that the problem continues to remain substantially unsolved.

2 Some well-known indices of poverty

Let $\mathbf{X}_n$ be the set of all non-decreasingly ordered non-negative n-vectors of income $\mathbf{x} = (x_1,...,x_i,\ldots,x_n)$, where x_i is the income of the ith poorest person in a community of n individuals, and $0 \leq x_i \leq x_{i+1}$, for all i = 1,…,n-1. Let N be the set of positive integers, and $\mathbf{X}$ the set $\cup_{n\in N}\mathbf{X}_n$. Let R be the real line, and S the positive real line. z (> 0) will be taken to represent the *poverty line*; and given a set N of n persons confronted by an income vector $\mathbf{x}$, the set Q of poor persons is defined by $Q(\mathbf{x};z) \equiv \{i\in N | x_i < z\}$. The cardinality of Q will be denoted by q. A *poverty index* is a mapping $P:\mathbf{X}\text{x}S \to R$ such that, for every $(\mathbf{x};z)$ in its domain, P specifies a unique real number which measures the extent of poverty associated with $(\mathbf{x};z)$. Throughout, we shall take it that n is given and fixed (so that all our poverty comparisons will be 'fixed population' poverty comparisons).

Here are some well-known poverty indices—defined for all $(\mathbf{x};z)\in\mathbf{X}\text{x}S$—which have been suggested in the literature.

The headcount ratio (H). The headcount ratio is the proportion of poor people in the population:

$$H(\mathbf{x};z) = q(\mathbf{x};z)/n(\mathbf{x}).$$

The income gap ratio (I). Let μ^P be the average income of the poor: $\mu^P(\mathbf{x};z) \equiv [1/q(\mathbf{x};z)]\Sigma_{i\in Q(x;z)}x_i$. Then, the income gap ratio I is defined as the proportionate shortfall of the average income of the poor from the poverty line:

$$I(\mathbf{x};z) = 1 - \mu^P(\mathbf{x};z)/z.$$

The per capita income gap ratio (R). R is the aggregate income deficit of the poor, expressed as a proportion of the total income that would be needed to raise them to the poverty line, and is given by the product of H and I:

$$R(\mathbf{x};z) = H(\mathbf{x};z).I(\mathbf{x};z).$$

Sen's Index of Poverty (P_S). Sen's poverty index (see Sen 1976a) is given by the expression:

$P_S(\mathbf{x};z) = [2/\{q(\mathbf{x};z)+1\}z]\Sigma_{i\in Q(\mathbf{x};z)}(z-x_i)(q(\mathbf{x};z)+1-i)$.

For ‘large’ values of q, we can write the Sen index, in an asymptotic approximation, as

$P_S(\mathbf{x};z) = H(\mathbf{x};z)[I(\mathbf{x};z) + \{1- I(\mathbf{x};z)\}.G(\mathbf{x};z)]$,

where G [$\equiv 1 + 1/q -(2/q^2\mu^P)\Sigma_{i\in Q}x_i(q+1-i)$] is the Gini coefficient of inequality in the distribution of poor incomes.

It may be noted that H violates the *monotonicity axiom*, which is the requirement that, other things equal, a diminution in a poor person’s income should cause measured poverty to rise; I and R satisfy the monotonicity axiom but violate the *weak transfer axiom*, which is the requirement that, other things equal, a transfer of income from a poor person to a richer poor person which keeps the latter poor should cause poverty to rise; and the Sen index satisfies both the monotonicity and the weak transfer axioms.

The Foster-Greer-Thorbecke (or FGT) P_α family of indices. Foster, Greer and Thorbecke (1984) proposed an entire class of poverty indices, parametrized by the quantity α (which is to be interpreted as a measure of ‘poverty aversion’), and given by:

$P_\alpha(\mathbf{x};z) = (1/n)\Sigma_{i\in Q(\mathbf{x};z)}\{(z-x_i)/z\}^\alpha, \alpha \geq 0$.

Some special cases of P_α are worth noting. P_0 is just the headcount ratio H, while P_1 is the per capita income gap ratio R. P_2 resembles the Sen index, but measures income inequality among the poor in terms of the squared coefficient of variation, C^2, rather than the Gini coefficient G: $P_2 = H[I^2 + (1-I)^2C^2]$. P_α satisfies monotonicity for $\alpha > 0$, and, additionally, transfer for $\alpha > 1$. For $\alpha > 2$, P_α satisfies a strengthened version of transfer, called ‘transfer sensitivity’, which makes the poverty index more sensitive in a specific way to income transfers at the lower than at the upper end of the income distribution of the poor. Indeed, as α keeps increasing, P_α becomes more and more distributionally sensitive until, in the limit, as α goes to infinity, P_α ranks distributions by the *maximin criterion*; i.e., according to the income level of the poorest individual. For all values of α, P_α is *decomposable*, that is, expressible as a population share-weighted sum of subgroup poverty levels. The FGT P_α family of poverty indices is perhaps the most widely employed set of poverty measures in use.

3 Poverty indices and moral intuition

In what follows, I shall present a set of elementary arithmetical examples which suggest that specific poverty judgments delivered by each of the poverty measures considered in Section 2 can militate against our moral intuition. We shall take the poverty line z to be given by z = rupees 100, and n (unless otherwise stated) to be 1 million.

First, consider the two n-person distributons $\mathbf{x}^1$ and $\mathbf{y}^1$, given, respectively, by $\mathbf{x}^1 = (99,...,99)$ and $\mathbf{y}^1 = (0,\ldots,0)$. Suppose we measure poverty by the headcount ratio. It is easily checked that $H(\mathbf{x}^1;z) = H(\mathbf{y}^1;z) = 1$. It is immediate that in terms of both the quantity and quality of deprivation, $\mathbf{x}^1$ should be regarded as being better than $\mathbf{y}^1$. The headcount ratio militates against this judgment by attending neither to quantity nor quality, but concentrating solely on the *numbers* in poverty: this is the well-known consequence of H's violation of the monotonicity axiom.

Next, suppose as before that n = 1 million. Let the n-vectors $\mathbf{x}^2$ and $\mathbf{y}^2$ be given, respectively, by $\mathbf{x}^2 = (0,99,\ldots,99)$ and $\mathbf{y}^2 = (0,\ldots,0)$. If we measure poverty by $P_{\alpha\to\infty}$, then the maximin criterion will come into play; and since the worst off individual in both distributions is equally badly off (with an income of 0), we must have: $P_{\alpha\to\infty}(\mathbf{x}^2;z) = P_{\alpha\to\infty}(\mathbf{y}^2;z)$, whereas, in terms of both the quantity and quality of deprivation, one would be inclined to judge that $\mathbf{x}^2$ is a better distribution than $\mathbf{y}^2$. As in the previous case, the poverty index under consideration ignores quantity and quality alike, concentrating, as it does, solely on a comparison of the worst off individual's fortunes: this is the well-known consequence of maximin's propensity to uphold a sort of 'dictatorship of the weakest'.

Now consider the n-dimensional distributions $\mathbf{x}^3 = (0,100,\ldots,100)$ and $\mathbf{y}^3 = (0,\ldots,0)$. If we measure poverty by the income gap ratio, then noting that the average income of the poor is 0 in both distributions, one must have: $I(\mathbf{x}^3;z) = I(\mathbf{y}^3;z) = 1$. It is hard to believe that $\mathbf{x}^3$, in which only one person out of a million is subjected to total deprivation, is not better than $\mathbf{y}^3$, in which all one million people out of a million are subjected to total deprivation. The index I attends only to quality, ignoring quantity altogether. Sen (1981: 33) cites as one of the 'damaging limitations' of I the fact that it '...pays no attention whatever to the number or proportion of ... people below the poverty line, concentrating only on the aggregate shortfall...'. Indeed, the problem shows up more acutely in the following example, where $\mathbf{x}^{3'} = (0,100,\ldots,100)$ and $\mathbf{y}^{3'} = (0.01,\ldots,0.01)$: $I(\mathbf{x}^{3'};z)$ [= 1] > $I(\mathbf{y}^{3'};z)$ [= 0.9999]. This corresponds to Parfit's (1984: 406) 'Two Hells' example:

> *The Two Hells*. In *Hell One*, the last generation consists of ten innocent people, who each suffer great agony for fifty years. The lives of these people are much worse than nothing. They would all kill themselves if they could. In *Hell Two*, the last generation consists not of ten but of ten million innocent people, who each suffer agony just as great for fifty years minus a day.

With an appropriate change of circumstantial detail, $\mathbf{x}^{3'}$ and $\mathbf{y}^{3'}$ can be seen to correspond to Hell 1 and Hell 2 respectively. In certifying $\mathbf{x}^{3'}$ to be poverty-wise worse than $\mathbf{y}^{3'}$, the index I is relying solely on what Parfit calls the 'Average Principle'; but, as he observes (1984: 406) 'When we consider these imagined Hells, we cannot plausibly appeal only to the … Average View'.

Perhaps the per capita income gap ratio R ($\equiv$ HI) will fare better? Let $\mathbf{x}^4$ and $\mathbf{y}^4$ be two one-million-vectors, with $\mathbf{x}^4 = (99.9999,\ldots,99.9999)$ and $\mathbf{y}^4 = (0,100,\ldots,100)$. Note that in neither case is there any inequality in the distribution of poor incomes, whence the Sen index of poverty, in these cases, will coincide with the per capita income gap ratio. It is easy to verify that $R(\mathbf{x}^4;z)$ [$= P_S(\mathbf{x}^4;z)$] $= R(\mathbf{y}^4;z)$ [$= P_S(\mathbf{y}^4;z)$] $= 10^{-6}$. The judgment delivered by R and P_S is scarcely convincing. Note that $\mathbf{y}^4$ can be derived from $\mathbf{x}^4$ by depriving person 1 of his entire income and distributing it equally among the remaining 999,999 individuals, so that each of these individuals is better off to the extent of rupees 0.0001. The tradeoff, rightly, appears to be absurd. What is at work is what Parfit calls the 'Total Principle'—the concern is entirely with quantity, and not at all with quality. And it is precisely this sort of obsession with total quantity which leads to Parfit's criticism of classical utilitarianism, in terms of its precipitating his Repugnant Conclusion. A poverty version of the Repugnant Conclusion would sound, loosely, something like the following: 'One must be poverty-wise morally indifferent between two situations, in one of which a very large number of people experience relatively low levels of deprivation (with incomes close to the poverty line), and in the other, a single person is made to bear the brunt of complete destitution—just as long as the total quantity of deprivation in the two situations is the same'.

The P_α indices, for finite values of α, fare no better. Take any positive, finite value of α, and suppose $n = 10^\alpha$. Suppose also, as before, that z = Rupees 100. Consider a pair of distributions $(\mathbf{x}^5, \mathbf{y}^5)$, with $\mathbf{x}^5 = (90,\ldots,90)$ and $\mathbf{y}^5 = (0,100,\ldots,100)$. The common tendency would be to judge $\mathbf{x}^5$ to be a poverty-wise better distribution than $\mathbf{y}^5$. But again, because of a complete commitment to quantity at the expense of quality, P_α will force indifference between the two distributions: it is easily verified that $P_\alpha(\mathbf{x}^5;z) = P_\alpha(\mathbf{y}^5;z) = 1/n$.

By considering pairs of distributions like $(\mathbf{x}^1,\mathbf{y}^1)$, $(\mathbf{x}^2,\mathbf{y}^2)$, $(\mathbf{x}^3,\mathbf{y}^3)$, $(\mathbf{x}^4,\mathbf{y}^4)$, and $(\mathbf{x}^5,\mathbf{y}^5)$, it has been shown that each of the poverty measures H, I, R, P_S, $P_{\alpha(\text{finite})}$, and $P_{\alpha\to\infty}$ can precipitate poverty comparisons which go against our moral intuition, and also against our considered judgment on the value of both quantity and quality in an overall appraisal of poverty. These indices, in specific instances, ignore both quantity and quality, or stress quality at the expense of quantity, or quantity at the expense of quality. The problem, as has been noted by Subramanian (1989: 245), is that '… a unique real number imposes an artificial moral equivalence among various outcomes with vastly differing ethical implications'. While other poverty indices proposed in the literature could be considered, or a more general statement of the problem could be essayed, there is enough on board, I believe, to suggest that there are difficulties with poverty aggregation which are deep, and endemic to the considerations which have conventionally inspired it. It would appear that the programme of poverty measurement, as it has evolved over time, and the programme of constructing a satisfactory theory of beneficence, have intersecting areas of both ambition and failure.

4 A way out?

What we have considered so far in this note is only the *aggregation* exercise in poverty measurement, having implicitly assumed that the *identification* problem—that of fixing the poverty line z—has been satisfactorily solved. In point of fact, though, the latter problem is an enormously difficult one to solve. Should the poverty line be seen as an absolute one or a relative one? What exactly does one mean by 'relative'? Should z be determined through a 'basic needs' approach? What are these basic needs, which commodities satisfy them, how much of each of these commodities is the 'right' amount? Should the poverty line be fixed with respect to nutritional requirements? Which nutrients, in what quantities, and with what allowance, if any, for intra- and inter-individual variations in requirements? And so on, and on. In dealing with these questions, one has the luxury of neither a reasonable absence of ambiguity at the conceptual level, nor a reasonable presence of consensual agreement on the ground.

At the level of logic, a possible way out of the aggregation problem considered earlier, is one which could also, and at the same time, address the identification problem. This constructive suggestion is offered here in a spirit of considerable tentativeness—for reasons which will be amplified in the following section. It entails no more than a direct application of Sen's (1976b) distributionally adjusted welfare measure to the poverty problem, along the following lines. Letting D stand for a measure of inequality in the distribution of poor incomes, the aggregation problem, it may be recalled, centers around identifying a 'satisfactory' way in which H, I and D can be combined. Where there is no inequality, the problem is one of finding a satisfactory way in which H and I can be combined. A widely accepted recourse has been to the multiplicative form HI, which corresponds to the 'quantity' component of deprivation (while I corresponds to its 'quality' component). For a fixed population, both H and I can vary, and it is variations in *both* which causes the problem. Logically, therefore, a natural way out which suggests itself is to get rid of either H or I, so that one is left to deal with only one dimension of poverty. Arising from this, one is tempted to ask: what if one simply fixed H once and for all, at say 50 per cent? Then, effectively, one only has to worry about the 'quality' component of deprivation (and, where there is inequality in the distribution of deprivation, about such inequality). For all practical purposes, then, the identification problem has also been solved: the poverty line z is simply the median income, call it M. Let μ^M be the average income of the 'poor', *now interpreted simply as the poorest 50 per cent of any population*; and let G^M be the Gini coefficient of inequality in the distribution of incomes among the poorest 50 per cent of the population. Sen's distributionally adjusted welfare measure, restricted to the poor, is the quantity $\mu^M(1- G^M)$. The *negative* of this, namely $\mu^M(G^M-1)$, it is being suggested here, be used as an index of 'illfare of the poor', or simply—if it does not entail too much of an abuse of language—as a poverty index, call it Π. Π, it can be checked, will lie in the interval $(-\infty,0]$.

It may be objected that the poor are always going to be with us according to the proposed measure. While this is true, it may also be held that this should not cause too much concern, so long as our purpose in measuring poverty is inspired only by the need to effect poverty *comparisons*. Additionally, if one treats it as a convention to view the poorest 50 per cent of any population as ‘the poor’, one does not have to any longer contend with the identification problem. Moreover, the index Π does not lend itself to any of the odd poverty judgments that were reviewed in Section 3 with respect to the pairs of distributions ($\mathbf{x}^1$,$\mathbf{y}^1$), ($\mathbf{x}^2$,$\mathbf{y}^2$), ($\mathbf{x}^3$,$\mathbf{y}^3$), ($\mathbf{x}^4$,$\mathbf{y}^4$), and ($\mathbf{x}^5$,$\mathbf{y}^5$). And finally, the index Π is closely related to the Sen index of poverty P_S—an unsurprising fact, considering that both indices were created by the same person. It is easily verified that when q is ‘large’, the Sen poverty index can be written as

$$P_S = H – H(\mu^P/z)(1\text{-}G).$$

If we take z to be the median income M, then H becomes 50 per cent for all populations under comparison, and since it is a constant across the board, it can be eliminated from further consideration; μ^P becomes μ^M, and G becomes G^M; and by removing z from the picture, we avoid normalizing with respect to the poverty line. What effectively remains is $\mu^M(G^M\text{-}1)$, or Π.

Briefly, Π would seem to have a number of advantages *vis-à-vis* many poverty indices in current use, and it also has the merit of simplicity. Does this extricate one from the nihilism of Section 3? Unfortunately not, as the discussion in Section 5 reveals.

5 Not a way out!

While the ‘resolution’ suggested above may have some surface appeal from a logical point of view, it falls foul of persuasiveness when judged by a simple sociological criterion, namely the view, as Sen (1983, 1985: 332) has convincingly argued, that ‘[t]here is … an irreducible absolutist core in the idea of poverty’. Specifically, it could be objected that calling the poorest 50 per cent of any population ‘the poor’ amounts to Humpty-Dumptyism (‘When *I* use a word, it means just what I choose it to mean—neither more nor less’). It could be objected, in particular, that the poorest 50 per cent in one population could include in their lot some, or many, who are leading lives well beyond what would be dictated by their ‘basic needs’; and that the poorest 50 per cent in another population could exclude from their lot some, or many, who are leading lives well within what would be dictated by their ‘basic needs’. This must be regarded as a significant criticism. A simple numerical example should help to highlight the import of the difficulty. Let us suppose that it is possible, in some ‘rough’ yet meaningful way, to suggest that a person with income less than some z* is ‘absolutely impoverished’: whether or not a person is in absolute poverty is frequently a matter of practical knowledge, irrespective of the difficulty of stipulating a cut-off level of income with any precision. Now consider two equidimensional income vectors **x** and **y**, and an income level $x < z^*$, such that, in vector **x**, every person has an income of x while in vector **y**, half of the population have an income of x and the other half an income of z*.

Our normative instinct would dictate that **x** is poverty-wise worse than **y**, since 100 per cent of the population in **x** are 'absolutely impoverished', while this is true for only 50 per cent of the population in **y**: other things equal, the numbers in absolute poverty do matter. Yet, and as can be immediately seen, the index Π will force indifference between the distributions **x** and **y**: Π(**x**) = Π(**y**) = (-)x.

The embarrassment to one's intuition can be even more acute. Imagine two one-million-vectors **x** and **y**, such that in **x**, every person has an income of 1,000 rupees, while in **y**, the poorest fourth have an income of zero each and the richest three-fourths have an income of a million rupees each. Assume that z* can be pitched in the region of rupees 100. Then, clearly, Π(**x**) = (-)1,000; and, noting that $\mu^M(\mathbf{y}) = 500{,}000$ and $G^M(\mathbf{y}) = 0.5$, it must be the case that Π(**y**) = (-)250,000. Briefly, the index Π suggests that there is vastly less poverty in **y** than in **x**, even though, from any reasonably absolutist perspective on poverty, there is no (absolutely) poor person in **x**, while there are a quarter of a million completely impoverished individuals in **y**!

In short, if one concedes that 'there is … an irreducible absolutist core in the idea of poverty' (and it is hard to see how one can fail to do so), then the way out explored in Section 4 is really no way out at all. We are now enabled to see how the identification and the aggregation exercises—which are often treated as independent problems in the measurement of poverty—are actually deeply intertwined. If we were free to divorce the identification of the poverty line from any 'absolutist' considerations, then a measure like Π might be a reasonably satisfactory way of addressing the aggregation problem. But, since we do not have this freedom, we are forced back into the aggregational entanglements reviewed in Section 3.

6 Concluding observations

This note has pointed to certain correspondences between the project of poverty measurement as initiated by Amartya Sen, and the project of coming up with a satisfactory theory of beneficence, designed to address questions in population ethics, as initiated by Derek Parfit. Similarities of motivation, content, and outcome have been noted; and while both projects, it could be argued, have met with a failure of sorts, the failure has also been instructive with respect to the complexities that must be encountered and dealt with in the kinds of projects under review. At a more constructive level, this note has also proposed a possible way out of the poverty measurement problems reviewed in it, by, as it happens, pressing Amartya Sen's own work on 'real national income'—done more or less concurrently with his early work on poverty measurement—to the rescue. However, closer scrutiny of this approach has also pointed in the direction of failure.

Poverty statistics are a key pointer to how well or badly an economy is faring. Description, evaluation, and policy intervention require that both the identification and the aggregation exercises be addressed with care. Both magnitudes and directions of change/difference in poverty levels are highly sensitive to how the poverty line is fixed and what index is used to measure poverty. Reddy and Pogge (2003) show that international comparisons of poverty, as well as assessments of global trends in poverty, are only as good as the care with which the conceptual underpinnings of poverty measurement are addressed. Subramanian (2005) is a similar comment on the validity of India's official time-series poverty statistics. The present note suggests that the problem is further compounded by conceptual ambiguities at the very heart of poverty measurement. At the risk of pronouncing a prescriptive banality, it must be concluded that poverty is too important a subject for the analyst to measure it without care, or to interpret her findings without a self-conscious awareness of the logical problems inherent in its measurement.

Appendix

Some difficulties of combining the quantity and quality of wellbeing in addressing questions in population ethics: a brief exposition

Parfit (1984) points to two principles, which he calls, respectively, the Total Principle and the Average Principle, which could be summoned in addressing questions relating to the comparison of wellbeing in alternative population scenarios. The Total Principle is concerned only with the *quantity* of wellbeing, while the Average Principle is concerned only with its *quality*. Neither by itself is a satisfactory guide to wellbeing comparisons; furthermore, problems arise when we attempt to combine the two in a theory of beneficence. The nature of these difficulties is briefly explicated in what follows.

We consider the Total Principle first, and draw on Dasgupta (1993) for the ensuing discussion. Consider a society with a population of size M in which each person experiences a level of welfare w^*. Let the aggregate welfare in this society be represented by $W(w^*,M)$. Consider another possible society, with a population of size M+1, in which again each person experiences a welfare level of w^*, so that aggregate welfare in this society is $W(w^*,M+1)$. With Dasgupta (op. cit.), let us consider the class of ethical theories which presume the existence of a unique w^* such that, for all $M \geq 0$, $W(w^*,M) = W(w^*,M+1)$: then, this w^* can be calibrated as the zero level of welfare. Utilitarianism, which belongs to the class of ethical theories just alluded to, judges the aggregate welfare of a society to be given by the sum of individual welfare levels in that society. In line with this, for a society of M persons each of whom experiences a positive level of welfare w~, utilitarianism will dictate that aggregate welfare of the society is given by $W(w\sim,M) = Mw\sim$. Let us add k individuals to this society, and let w_k be the welfare level of each individual in this enhanced population. The utilitarian social welfare function will now pronounce that $W(w_k,M+k) = (M+k)w_k$, and will judge the two regimes to be welfare-equivalent if $W(w_k,M+k) = W(w\sim,M)$, or, equivalently, if $w_k = [M/(M+k)]w\sim$. Notice now that as $k \rightarrow \infty$, $w_k \rightarrow 0$: this is the Repugnant Conclusion—the conclusion, yielded by utilitarianism, which supports the ethical unexceptionableness of having indefinitely large populations whose members experience arbitrarily small positive levels of welfare. This is the consequence of an exclusive concern with the *quantity* of wellbeing.

Next, we consider the Average Principle. If w_i is the welfare level of the ith person in a society of M individuals, and if $m_M \equiv (1/M)\Sigma_{i=1}^{M} w_i$ is the average level of welfare of the society, then the Average Principle will certify social welfare to be given by $W(\{w_i\},M) = m_M$. Consider a two-person society (say, Adam and Eve) each of whom enjoys a high level of welfare of w^+; by the Average Principle, $W(\{w^+,w^+\},2) = w^+$. Consider an M-person society, with $M >> 2$, such that each person in this society (including Adam and Eve) enjoys a welfare level of $w^+-\varepsilon$, where ε is some arbitrarily small positive number. By the

Average Principle, $W(\{w^+-\varepsilon,\ldots, w^+-\varepsilon\},M) = w^+-\varepsilon$. The Average Principle asserts that it is best that Adam and Eve be the only persons to exist on earth: even a tiny reduction in their welfare level will not be seen to be compensated by the vastly larger numbers of people that are enabled to enjoy a high level of welfare.

The Total Principle offends our intuition because it upholds quantity at the complete expense of quality, and the Average Principle offends our intuition because it upholds quality at the complete expense of quantity. Parfit is thus led to suggest that *both* quantity and quality have a rightful place in a theory of beneficence. But combining both also leads to difficulties, as demonstrated below. Recall that the Total Principle permits average welfare to keep falling (for as long as it is positive) if the population increases by a sufficiently large amount to maintain the total quantity of welfare. We may wish to restrict the role of quantity by postulating some finite population size M* such that, for all populations M less than M*, aggregate welfare is Mm_M, and for all populations M greater than or equal to M*, aggregate welfare is M^*m_M. In this view, beyond a certain population size, the addition of lives with positive welfare ceases to automatically add to moral value. This, however, does not apply to lives with negative value, that is, to lives which are 'bad' or 'unhappy'—the addition of a 'bad' life always increases moral disvalue, without limit. Thus, if a proportion g of a population of size M have 'good' lives with a welfare level of w each, and a proportion 1-g have 'bad' lives with an *illfare* level of d each, then we may postulate a formula of aggregate net welfare given by

$$\begin{aligned} W(g,w,d,M) &= M[gw - (1-g)d] \text{ for all } M < M^*; \\ &= M^*gw - M(1-g)d \text{ for all } M \geq M^*. \end{aligned}$$

Consider three population scenarios A, B and C, where A will be taken to denote the present, and B and C are set in the future. Suppose the present population to be M*/2. Then the aggregate net welfare level under situation A is given by $W_A = (M^*/2)[gw - (1-g)d]$. For specificity, set g = 0.9 and d = 5w. Then, making the appropriate substitutions in the preceding equation, we have:

(1)$W_A = 0.2M^*w$.

In situation B, the future consists of two centuries in each of which the population is M*: every other feature of the world is the same as in situation A. It is then easy to see that in each of these centuries, the aggregate net welfare is given by $W_B = M^*[gw - (1-g)d] = 2W_A$, whence (in view of (1)), in each of these centuries:

(2) $W_B = 0.4M^*w$.

The future, as represented by situation B, is clearly better than the present, as represented by situation A. Finally, situation C will be taken to be identical to situation B, save that in situation C the future consists of only one century: the aggregate net welfare level in situation C will then be given by $W_C = M^*gw - 2M^*(1\text{-}g)w$, which, given the assumed parameter values, reduces to

(3) $W_C = - 0.1M^*w$.

From (3) we notice that aggregate net welfare in situation C is actually *negative*: it would be better if no-one existed in situation C. From (1), (2) and (3), we must conclude that the future as represented by B is better than the present, while the future as represented by C is worse than the present and, indeed, intrinsically bad. Yet, in all relevant details, the two futures are identical—with respect to the size of the total population, and with respect to the distribution of 'good' and 'bad' lives. This conclusion to which we are driven, in an attempt to meaningfully combine notions of both quantity and quality in a theory of beneficence, is what Parfit calls the 'Absurd Conclusion'.

To summarize, in isolation or together, we encounter problems of coherence in accommodating considerations of quantity and quality in an assessment of human wellbeing.

References

Dasgupta, P.S. (1993): *An Enquiry in Wellbeing and Destitution*, Clarendon Press: Oxford.

Foster, J., J. Greer, and E. Thorbecke (1984): 'A Class of Decomposable Poverty Measures', *Econometrica*, 52(3): 761-6.

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

Reddy, S.G., and T.W. Pogge (2003): 'How *Not* To Count The Poor', forthcoming in S. Anand and J. Stiglitz (eds): *Measuring Global Poverty,* Oxford University Press: Oxford. Also available at www.socialanalysis.org.

Sen, A.K. (1976a). 'Poverty: An Ordinal Approach to Measurement', *Econometrica*, 44(2): 219-31. Reprinted in S. Subramanian (ed.) (1997): *Measurement of Inequality and Poverty*, Oxford University Press: Delhi.

Sen, A.K. (1976b): 'Real National Income', *Review of Economic Studies*, 43.

Sen, A.K. (1981): *Poverty and Famines: An Essay on Entitlement and Deprivation*, Clarendon Press: Oxford.

Sen, A.K. (1983): 'Poor, Relatively Speaking', *Oxford Economic Papers*, 35. Reprinted in A.K. Sen (1985): *Resources, Values and Development*, Oxford University Press: Delhi.

Subramanian, S. (1989): 'Poverty-Minimization in the Light of "Life-Boat Ethics"', *Journal of Quantitative Economics*, 5.

Subramanian, S. (2005): 'Unravelling a Conceptual Muddle: India's Poverty Statistics in the Light of Basic Demand Theory', *Economic and Political Weekly*, January 1: 55-66.

Watts, H. (1968): 'An Economic Definition of Poverty', in D.P. Moynihan (ed.), *On Understanding Poverty*, Basic Books: New York.