Discontinuous Losses from Poverty, Generalized P_α Measures, and Optimal Transfers to the Poor

Gary Fields

Department of Economics, Cornell University

François Bourguignon

DELTA-ENS

Journal of Public Economics (1997), 63, 155-175

An earlier version of this paper was presented to the conference on 'Poverty and Human Resources in Developing Countries and their Implications for Policy', Institute for International Economics, Stockholm, Sweden, August 1992. We thank J. Foster and an anonymous referee for useful comments.

Abstract

This paper examines the distributional properties of poverty measures which are discontinuous at the poverty line. It is shown that among all the additive poverty measures, only those measures with some discontinuous jump at the poverty line are such that it is optimal to allocate a given antipoverty budget either to the richest of the poor, or to the poorest of the poor, or to both. A special class of such poverty measures is an extension of the well-known P_α, the properties of which are investigated.

Keywords: Poverty index; Antipoverty policy

1. Introduction

Poverty measures can be interpreted as gauging the social welfare losses when persons have low incomes.[1] Two distinct aspects of such losses may be identified: the loss from being poor and the loss from being poorer. The loss from being poor arises whenever a person is too poor to acquire the poverty-level basket of goods. The loss from being poorer reflects the judgment that poverty is more severe the poorer a poor person is. The commonly-used poverty headcount ratio deals only with the first of these, while the P_α class of poverty measures (Foster et al., 1984) deals only with the second.

In Section 2 of this paper we discuss the loss from being poor and seek to justify why the loss-from-poverty function might be discontinuous at the poverty line. The social welfare justification for this discontinuity harkens back to the linen shirt' argument of Adam Smith and to the 'capability' concept advocated by Amartya Sen.

Section 3 amalgamates the notions of the loss from being poor and the loss from being poorer into a single loss-from-poverty function. We specify a general class of such functions and analyze the poverty-minimizing allocation of an antipoverty budget for this class. In previous work (Bourguignon and Fields, 1990) we showed that the optimal allocation using various poverty measures could be of three different types: (i) allocating all the budget to the poorest of the poor, whose needs are greatest; (ii) allocating all the budget to the richest of the poor, so that as many people as possible can escape poverty; or (iii) making a mixed allocation with part of the budget going to the poorest of the poor and part to the richest of the poor. We show in the present paper that, within a fairly general class of poverty measures, only those who are

---

[1] For rotational convenience, throughout this paper 'person' is the recipient unit and 'income' is the basis on which poverty and standard of living are based.

discontinuous at the poverty line—and thus exhibit some fixed loss from poverty—may be consistent with the third type of allocation, and we analyze the characteristics of that allocation.

Section 4 then develops a new extension of the P_α index which is a special case of the general class analyzed in Section 3. This new $P_{\alpha,\delta}$ class retains all the axioms and derivative properties of the P_α index, while also combining with them the insight reflected in the headcount ratio on the loss from being poor. This fixed loss from poverty causes the 'generalized $P_{\delta\alpha}$' class to exhibit allocative properties which may be quite different from those implied by the standard P_α measure, and which may be more appealing for practical use in public finance.

Remarks and extensions appear in Section 5. Section 6 concludes.

2. Justification for a discontinuous social loss-from-poverty function at the poverty line

A long-standing debate in the poverty literature is whether poverty is a discrete or a continuous phenomenon.[2] Since this paper investigates the implications of regarding the social loss from poverty as discontinuous at the poverty line, it is important to consider on what fundamental social welfare grounds such a discontinuity may be justified.

Let us stress first that the very fact that the social observer, or planner, may be interested in poverty rather than, more generally, in the income distribution throughout society reveals that being in poverty is qualitatively different from having a low income in the observer's social valuation of income. In particular, individual incomes matter below some line z and they do not matter above it. This implies that the social loss from poverty is unaffected by income changes among the non-poor; what matters are redistributions to the poor from the non-poor and

[2] See, for instance, Atkinson (1987), Lewis and Ulph (1988), and Lipton and Ravallion (1995).

redistributions among the poor. In the additively separable case with which we work in this paper, this means that the social loss from poverty is zero for incomes above the poverty line and positive for incomes below it.

The preceding argument says that the practice of focusing on the poor segment of the population rather than on the whole population is not in contradiction with assuming discontinuity of the social loss from poverty. But it is also consistent with continuous loss-from-poverty functions. This is in fact the presentation adopted in the recent literature which applies to poverty the same ‘social dominance' concepts as those developed in the field of income distribution comparisons (e.g. Atkinson, 1987; Foster and Shorrocks, 1988; Ravallion, 1994).

The arguments in favor of a discontinuity are based on the idea that a minimum income is needed for an individual to perform ‘normally’ in a given environment and society. Below that income level some basic function of physical or social life cannot be fulfilled and the individual is somehow excluded from society, either in a physical sense (e.g. the long-run effects of an insufficient diet) or in a social sense (e.g. the ostracism against individuals not wearing the proper clothes, or having the proper consumption behavior).

This idea goes back to Adam Smith. Amartya Sen is fond of quoting Smith as follows: “In the present time, through the greater part of Europe, a creditable day-labourer would be ashamed to appear in public without a linen shirt.” Sen himself adds: “On the space of the capabilities themselves - the direct constituent of the standard of living - escape from poverty has an absolute requirement, to wit, avoidance of this type of shame. Not so much having equal shame as others, but just not being ashamed, absolutely” (Sen, 1983, p. 335).

As we interpret the linen shirt argument, it has two powerful implications. First, an individual too poor to be able to buy a linen shirt suffers shame absolutely, i.e. either he has the means to buy the linen shirt and does so and is proud, or else he lacks the means to buy the shirt and suffers shame. Secondly, the shame he suffers is discrete-he suffers a full amount of shame even if he is only epsilon short of being able to buy the shirt.

It is in fact difficult to find in the literature a formal justification of the discontinuity analyzed in the present paper, even in the numerous writings by Sen. The reason may be that the linen shirt argument, or the capability approach to welfare, or the 'basic needs' concept, all lead, to a large extent, to a critique of the standard utilitarian-like social welfare approach to distributive issues. Our point here is precisely that, starting from the utilitarian framework, a step in the right direction pointed out by these remarks may be to introduce a discontinuity in the conventional social valuation of income, or a ‘fixed loss’ from poverty which arises in addition to the income-dependent ‘variable loss’ from poverty.

A good illustration of all the preceding points is provided by the familiar empirical Linear Expenditure System (LES) used in studies of household consumption behavior and their welfare implications. The LES consumption model is based on the Stone-Geary utility function:

$$u(x_1, x_2, \dots, x_n) = \sum_i \beta_i \log(x_i - c_i)$$

where x_i is the quantity consumed of good $i(= 1,2, \dots, n), \beta_i$ are positive parameters and c_i can be interpreted as some ‘minimum’ consumption of good i. When estimating and then using this model in the analysis of distributive issues, it must be assumed that all individuals have a consumption budget above the minimum, R_0, required to buy the ‘minimum’ bundle $(c_1, c_2, \dots, c_n)$, i.e. if p_i is the price of good i, then $R_0 = \Sigma p_i c_i$. What should be done if this assumption proves to be unrealistic? What must be done with the people whose budgets fall below that limit?[3] They may be considered precisely as poor. However, their consumption budget or their income clearly cannot be used as an index of their welfare in the same way as for households above the limit R_0 since they cannot have preferences represented by the LES model. If they cannot be compared with people who receive R_0 or slightly more then there is no reason to assume that their welfare level is the same as those people's when their own income tends toward R_0 from below. This is where we believe that some discontinuity must be introduced, although the extent of that discontinuity, i.e. the fixed social loss from poverty, must be set by the observer.

It is on that basis that we now consider a new class of poverty measures in which such a discontinuity is explicitly taken into account.

3. The general class of 'poverty line discontinuous' poverty measures

Let us consider a distribution of incomes among a population of N individual consumption units defined as $(y_1, y_2, \dots, y_N)$ where, without loss of generality $y_1 < y_2 < \cdots < y_N$. Let $(y_1, y_2, \dots, y_q)$ be the restriction of the previous distribution to the 'poor segment' of the population, in other words, q is the number of poor and we have that $y_1 < y_2 < \cdots < y_q < z < y_{q+1}$, where z is the poverty line. We consider in this paper the general class of additively separable poverty measures, the form of which is given by

$$P_L(y_1, y_2, \dots, y_N) = \frac{1}{N}\sum_{i=1}^{N} L(y_i)$$

[3] Nothing prevents this from happening at the estimation stage. Most authors using this type of model now indicate the percentage of households in this situation in their sample.

where $L(\cdot)$ is continuous, differentiable, non-increasing and (strictly) convex over the interval $[0, z[$, and $L(y) = 0$ over the interval $[z, \infty[$. It is through the latter restriction that the poverty measure $P_L(\cdot)$ is defined only on the poor segment of the population, and independent of the distribution of income among the non-poor. With that restriction, the function $L(\cdot)$ is to be interpreted as the social loss from the ith individual's poverty discussed in the preceding section, and $P_L(\cdot)$ as the mean social loss. One can check that members of the $P_L(\cdot)$ class defined by (1) satisfy monotonicity, subgroup consistency and decomposability.[4]

That $L(\cdot)$ is non-increasing guarantees that poverty cannot go up when the income of a poor person increases, whereas convexity implies that an increase of a given size in the income of a poor person reduces poverty by more the poorer the income recipient is.

The family of measures given by (1) includes the poverty gap, the popular P_α measures, and what Hagenaars (1987) called the 'Dalton class' (which includes, inter alia, the Watts index). The poverty gap obtains with $L(x) = (z - \mu)/z$, where μ is the mean income of the poor, P_α measures correspond to $L(x) = [(z - x)/z]^\alpha$, and the Dalton class to $L(x) = [u(z) - u(x)]/u(z)$, where u is increasing and concave.[5] Among the poverty measures that do not satisfy the preceding properties are the headcount ratio, the Sen index—and other measures derived from it—and the first family of indices proposed by Clark et al. (1981). As shown in Fig. 1(a), the implicit loss function in the headcount ratio is not decreasing and not strictly convex, whereas the inclusion of the Gini coefficient in the Sen index is equivalent to allowing the loss from poverty to depend not only on one's income but also on that of the other poor people. The reason is that unlike (1), the Sen index is not additively separable (see Section 5 below). Note though

[4] For definitions and discussions of these concepts, see Sen (1973), Foster et al. (1984), and Foster and Shorrocks (1991).
[5] It is well known that the poverty gap is P_α with $\alpha = 1$. Note also that only P_α. with $\alpha > 1$, are relevant here because of the strict convexity requirement in (1).

that this requirement in (1) may be somewhat attenuated by taking some monotonic transformation on the RHS of (1). Most of what follows actually applies to that wider class of 'weakly additive' poverty measures which, except for the convexity and differentiability requirement on the functions $L(\cdot)$, coincides with the family of 'subgroup consistent poverty indices' defined in Foster and Shorrocks (1991).

Insert Figure 1 Here

Following the argument in the preceding section, an important subclass of poverty measures which we emphasize in this paper is made up of functions of type (1), where $L(y)$ is discontinuous at the poverty line z with $L(z^{-}) = \delta$ and $L(z) = 0$, and $L(z^{-})$ is the limit of $L(y)$ when y tends towards z from below. We shall call this subclass of (1) the 'poverty-line-discontinuous' (PLD) class of measures—see Fig. 1(c).

We now investigate the properties of the general class of poverty measures (1) and that of the PLD subclass. We begin by investigating the optimal allocation of a given anti-poverty budget B, i.e. the allocation leading to the least poverty, as defined by a measure of type (1) or PLD.[6] As analyzed in a previous paper (Bourguignon and Fields, 1990), that optimal allocation may be of different types. It may be 'p-type', with all the budget B going to the poorest of the poor, or 'r-type' with all the budget being used to lift the richest of the poor out of poverty, i.e. to give them a final income equal to z. In some instances, it may also be 'mixed-type', i.e. a

[6] We only consider in this section the case of perfect targeting, i.e. the case where all transfers may be chanelled to specific individuals in the poor population without leakages. Imperfect targeting situations are considered in Section 5.

mixture of p-type and r-type. We shall show that in the family (1) of poverty measures, only the PLD subclass may lead to optimal allocations of the mixed-type.

Formally, the optimal allocation program is defined by

$$\min P_L(y_1 + t_1, y_2 + t_2, \dots, y_N + t_N) | t_1 + t_2 + \dots + t_N \leq B, \text{ with } t_i \geq 0, \forall i$$

Here, t_i is the income transfer made to individual i. It is required to be non-negative for all individuals, so that we do not really consider the way the antipoverty budget B is financed. This may be through taxes on all individuals, in which case y_i, is to be interpreted as income net of taxes. But B may also be some exogenous transfer made to the population, e.g. foreign aid. It may also be noted that, if P_L belongs to the family of poverty measures (1), then optimal transfers to the non-poor in (2) are zero, so that B may actually be financed as a tax on the non-poor. In any case, we shall take it for granted in what follows that t_i, is zero for $i > q$ and we shall consider only the determination of optimal transfers to the poor.

Our interest lies in the way the optimal transfer vector t^* depends on the function $L(\cdot)$ that defines the poverty measure $P_L(\cdot)$. It is easily shown, for instance, that with P_α and $\alpha > 1$, it is optimal to transfer the entire antipoverty budget to the poorest of the poor. The transfer thus is ‘p-type’ with

$$t_i^* = a - y_i \,; \qquad \text{for } i = 1,2, \dots, j$$

where a and j are given by

$$y_{j+1} > a > y_i; \qquad \sum_{i=1}^{i} (a - y_1) = B$$

Of course, this definition is valid only for a budget B that does not permit all the poor to be lifted out of poverty.

If we now consider the headcount ratio—which, as already noted, does not belong to the family of measures defined by (1)—the optimal transfer would be r-type and defined by

$$t_i^* = (z - y_i) \quad \text{for } i = q, q-1, q-2, \dots, q-j$$

where j is given by

$$\sum_{i=q-j}^{q} (z - y_i) = B$$

It must be noted, though, that this argument is valid only if B is exactly equal to that amount necessary to lift out of poverty an integer number of poor. In the case, for instance, where

$$\sum_{i=q-j+1}^{q} (z - y_i) = b_{j-1} < B < b_j = \sum_{i=q-j}^{q} (z - y_i)$$

only $j - 1$ poor can be lifted out of poverty, whereas there is some 'remainder' to be allocated. In the present case, the remainder is $R = B - b_{j-t}$. With the headcount ratio or the poverty gap, it does not matter how that remainder is allocated, but in the case of strictly convex loss functions in $[0, z[$, it would clearly be allocated to the poorest of the poor. We shall ignore that small imprecision in what follows by adopting a weak definition of the r-type transfer.

Definition of p-type, r-type and mixed-type transfers. (a) A transfer is *pure p-type* if it is entirely spent on the poorest of the poor so as to raise as many of them as possible to a common income above their initial one. (b) A transfer is *p-type* if some but not necessarily all of it is spent on the poorest of the poor, (c) A transfer is *pure r-type* if it is used to lift as many of the richest of the poor out of poverty as possible, (d) A transfer is *r-type* if some but not necessarily all of it is spent to lift the richest of the poor out of poverty, (e) A transfer is *mixed-type* if some but not all of it is spent to raise the incomes of the poorest of the poor and some but not as much as possible of it is spent to lift the richest of the poor out of poverty; that is, a mixed transfer consists of both p-type and r-type transfers.

Given these definitions, the first question we ask is: Under what conditions on the function $L(\cdot)$ might a mixed transfer be optimal for some antipoverty budget B? The answer to that question is given in the following three propositions.

Proposition 1. Among poverty measures of type (1) only PLD measures may lead to an optimal allocation of an antipoverty budget B of mixed-type.

Proof. Consider an optimal mixed type transfer. The poor fall into three categories: (i) those receiving a strictly positive transfer but still remaining below the poverty line (I); (ii) those who receive a strictly positive transfer which just enables them to escape poverty (J);[7] and (iii) those who do not receive any transfer at all (K). For such a transfer to be optimal, it must be the case that reducing the transfer to any poor $i \in J$ by any (positive) arbitrary small amount β and increasing by the same amount that of another poor $m \in I$ increases overall poverty. The change in the poverty measure corresponding to that transfer is given by

$$\Delta P_L = -L(z) + L(z-\beta) + \beta L'(y_m + t_m)$$

where the first two terms on the RHS correspond to individual i coming back into the pool of poor. Since the loss function is zero at z for all measures satisfying (1), optimality requires that

$$L(z-\beta)|\beta > -L'(y_m + t_m)$$

The differentiability of $L(\cdot)$ on $[0, z[$ implies that

[7] Since there is no gain in $P_L(\cdot)$ from moving.individuals strictly above the poverty line, individuals in (J) must end up exactly at the poverty line, z.

$$\lim_{\beta \to 0} L(z - \beta) = L(z^-) - \beta . L'(z^-)$$

where $L(z^-)$ and $L'(z^-)$ are, respectively, the limits of the loss function $L(y)$ and its derivative when y tends towards the poverty line z from below. After dividing by β, the inequality above becomes

$$\lim_{\beta \to 0} L(z^-) | \beta > L'(z^-) - L'(y_m + t_m)$$

But the RHS is strictly positive if $L(y)$ is strictly convex. Hence $L'(z^-)$ must be strictly positive if the optimal allocation is of the mixed-type. Q.E.D.

Now that we have shown that only PLD measures can lead to mixed transfers, we investigate under what conditions they actually do and in addition, when the optimal transfer is pure-r or pure-p type. We consider first the case of income distributions which are continuous in a neighborhood of the poverty line z. In other words, we assume for now that there is an individual with an income infinitely close to (and below) z.

Proposition 2. For all PLD poverty measures and all income distributions continuous in the neighborhood of z*, there exists a critical level* B*, of the antipoverty budget such that the optimal allocation is (a) pure r-type if* $B \leq B_r$ *and (b) mixed-type if* $B > B_r$*. The critical level* B_r *is defined in Eq. (4) below.*

Proof. By assumption, $L(\cdot)$ is decreasing, strictly convex and tends towards $\delta > 0$ as y tends towards z from below. Let us first show that, under the assumption of a continuous income distribution, the first cents of any antipoverty budget ($\in$) must be allocated in an r-type manner. Let $\in = z - y_q$ which, by continuity, is infinitely small. Without very much loss of generality, let us assume that $\in < (y_2 - y_1)$.[8] Under these conditions, e should be allocated to the poorest of the poor if the following condition is satisfied:

$$\in \cdot L'(y_1) < -L(y_q)(< -\delta)$$

This condition requires that the fall in poverty as a result of increasing the income of the poorest individual by ∈ be larger (in absolute value) than that as a result of lifting the richest individual out of poverty. However, because the richest poor person is assumed to be infinitely close to the poverty line and because there is a non-infinitesimal gain in welfare when that person escapes poverty, this condition cannot hold. By continuity, this shows that for an antipoverty budget small enough, the optimal transfer is necessarily pure r-type.

Until what point is this true? Assume that some amount has already been spent on r-type transfers so that there no poor person between the income level y^0 and the poverty line z. Keeping with an r-type strategy implies spending the next $z - y^0$ dollars on the richest of the poor, whereas making a p-type transfer would involve spending that amount on the poorest of the poor. The drop in poverty obtained by using those $z - y^0$ dollars for an r-type transfer is simply $L(y^0)$, whereas that obtained by using these dollars for a p-type transfer is given as follows. First, raise the poorest individual (whose income is y_1) to the next individual, then raise those two individuals to the next individual, and so on until the budget is exhausted. Denote the number of individuals whose incomes are so raised by i_0 and the amount to which their incomes are raised by y_0. Then the ith individual's loss from poverty is reduced by $L(y_i) - L(y_0)$, and overall poverty is reduced by:

$$M(y^0) = \sum_{i=1}^{i_0} [L(y_i) - L(y_0)]$$

[8] If this were not the case, then the allocation of € among the poorest persons should be taken into account.

$$\text{where} \quad y_0 \in [y_{i_0}, y_{i_0+1}]; \quad \sum_{i=1}^{i_0} (y_0 - y_i) = z - y^0$$

Note that, unlike in the proof of Proposition 1, the marginal loss function $L'(\cdot)$ is no longer what matters. This is because many people at the top of the distribution of income among the poor have already been lifted out of poverty and the amount $z - y^0$ to be allocated to the poorest is not infinitesimally small.

The income threshold at which it is no longer optimal to transfer to the richest of the poor is given by the equality between $L(y^0)$ and $M(y^0)$. We denote this threshold by y^*, the number of poor individuals who receive r-type transfers by j^*, and the number of poor individuals who would benefit from p-type transfers by i^*. The critical budget level B_r such that the optimal budget allocation switches from a pure r-type to a mixed-type is defined with reference to a threshold income y^* given by

$$L(y^*) = \sum_{i=1}^{i^*} [L(y_i) - L(y_0)]$$

$$\text{where} \quad y_0 \in [y_{i^*}, y_{i^*+1}]; \quad \sum_{i=1}^{i^*} (y_0 - y_i) = z - y^*$$

through the relationship

$$B_r = \sum_{i=j^*}^{q} (z - y_i), \quad where\ y_{j^*-1} < y^* < y_j$$

It remains to show whether $L(y^0)$ and $M(\quad ^0)$ can in fact be equated. The necessary analysis is performed with the aid of Fig. 2.

As noted above, the $L(y^0)$ function is the welfare gain from using an antipoverty budget of $z - y^0$ dollars to make an r-type transfer to a person with income y^0. This is simply the $L(y)$

function, i.e. it is decreasing, strictly convex, and tends towards $\delta > 0$ as y tends towards z from below. The $M(y^0)$ function is the welfare gain from using an antipoverty budget of $z - y^0$ dollars to make a p-type transfer to the poorest of the poor. We know that

$$\lim_{y^0 \to z} M(y^0) = 0$$

because at z, no money is available to give to the poorest of the poor. As y° falls, M(y") increases, because more and more money is available to give to the poor. The concavity of M(y°) can be proved in the case of a continuous distribution from the fact that M'(y') = L‘(y,,), where y,, is defined in (3), and clearly is a decreasing function of y°. Finally, the L(y") curve is always flatter than the M(y°) curve, i.e. |T’(y0)| < |A/'(y°)|, Vy"; this follows from the fact that the gain from a dollar going to the richest of the poor while he is still poor is less than the gain from an extra dollar going to the poorest of the poor.