



Queen's Economics Department Working Paper No. 822

Needs and Targeting

Michael Michael Keen

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

1-1991

NEEDS AND TARGETTING

by

Michael Keen*

Discussion Paper No. 822

First version: October 1989

This version: January 1991

* University of Essex, England, and Queen's University, Canada. I am grateful for their suggestions and comments to two referees, an Associate Editor, Tony Shorrocks and seminar participants at CORE, the London School of Economics, University College London and the University of Warwick. Errors and views are my own.

Abstract: This paper shows that when - in the manner of the recent targeting literature - the resources available for poverty relief are allocated across heterogeneous groups so as to minimise a 'well-behaved' index of aggregate poverty, the optimal response to an increase in the needs of some group may be to reduce the resources allocated to it. Necessary and sufficient conditions for optimally-targeted benefits to be inversely related to needs are established under two polar forms of poverty alleviation strategy (pure contingency and strict means-testing) and the wider methodological implications of this possibility discussed.

JEL Classification: 0243, 915

Keywords: Poverty, social security, needs

1. Introduction

The effective targeting on those who need them most of the resources available for the relief of poverty is a principal concern in the design of poverty alleviation programmes, and one which - in developed and developing countries alike - has been to the fore in policy discussion and evaluation over the past decade or so. The 1985 social security reforms in the UK, for example, were motivated largely and explicitly as an attempt to improve the targeting of limited resources, and similar notions also feature prominently in the 1989 World Development Report on Poverty. This practical emphasis has been matched by the emergence of a literature in which the targeting problem is formalised as one of minimising some explicit index of poverty subject to a revenue constraint: see, for instance, Kanbur (1987), Besley and Kanbur (1988), Besley (1990) and, for an empirical application, Ravallion and Chao (1989). Amongst the aspects of targeting that merit and have received particular attention is the optimal allocation of resources across heterogeneous groups: between one parent families and pensioner couples, for instance, or across rural workers in different regions. In practice, an important aspect of this heterogeneity will typically be a difference in the needs of these client groups, naturally characterised in the framework of the targeting literature as a difference in their poverty line incomes. Most income maintenance programmes incorporate relativities intended to reflect such differences in need, and indeed the selection of such scales remains a matter of considerable agonising. Yet the formal targeting literature has entirely ignored the issues raised by differences in needs: groups are treated as heterogeneous in terms, for instance, of the within-group distribution of income, but are almost invariably assumed to have the same poverty lines.¹ This instantly precludes any consideration of the relationship between the resources optimally targeted towards a group and the level of its needs.

The purpose of this paper is to argue that the relationship between needs and benefits that emerges from the targeting framework is far from

¹ Beath, Lewis and Ulph (1988) and Kanbur and Keen (1989) are exceptions, but neither develops the implications of differential poverty lines in any depth.

straightforward and, moreover, is conceptually troubling. In particular, it may be that optimal targeting requires that an increase in the needs of some group be met by a reduction in the resources allocated to it.² This possibility will be referred to - for want of anything handier - as the paradox³ of targeting. Its implications are striking, and perhaps perturbing. If the government behaves in the manner envisaged in the targeting literature, lobby groups may be well-advised, for instance, to press for a reduction in the poverty lines of their client groups.

It is useful to begin with a simple example. Consider a minister of state for social security who has a budget of £50 to spend on poverty relief, and who seeks in doing so to minimise the number of post-transfer poor; her minimand, that is, is simply the headcount measure of poverty. There are two client groups: group A consists of two individuals, each with pre-transfer incomes of £81; group B consists of a single individual with original income of £61. In this case the minister can afford to eliminate poverty altogether in either group A (by paying its members a benefit of £20) or group B (by a benefit of £40) but cannot afford to do both. The first of these strategies leaves a smaller proportion of the population in poverty (one-third as against two-thirds), and so will be chosen. But now suppose that for some reason - whether a change in political perceptions or a change in objective circumstances - the poverty line of group A rises to £110, that of group B being unchanged. Then it becomes impossible for the minister to take those in A out of poverty. The best she can now do is to concentrate resources on the elimination of poverty in B; and in the process, of course, benefits to those in the now-needier group A are cut.

This example establishes the logical possibility of the paradox. It also points to the consideration that underlies it: needy groups are expensive to help, so that one effect of an increase in needs is akin to a price increase, pushing towards a retargeting of resources in favour of groups

² This possibility is hinted at in Glewwe (1988).

³ The Oxford English Dictionary distinguishes several definitions of a 'paradox'. That intended here is the 1569 version: "A statement seemingly contradictory or absurd, though possibly well-founded or essentially true."

whose poverty is relatively cheap to alleviate. But the example is an extremely special one, and the intuition that it yields correspondingly blunt. Might the paradox still arise, for example, if the targeting problem is formulated in terms of poverty measures which - unlike the headcount ratio - are sensitive to the distribution of post-transfer income amongst the poor? Is it ruled out by, or even inherent in, particular poverty measures advocated in the literature and commonly used in practice? How does the shape of the pre-benefit income distribution affect the likelihood of its occurrence? How do the conditions for the paradox to occur vary with the nature of the poverty alleviation strategy that the government pursues?

Section 2 of the paper develops a simple framework for the analysis of poverty-minimising targeting across heterogeneous groups. The relationship between optimal benefits and needs is then examined under two polar assumptions on the form of the poverty alleviation strategy: the provision of purely contingent benefits (Section 3) and strict means-testing (Section 4). The results⁴ are summarised briefly in Section 5, which also takes up the central methodological question that they raise: Is the paradox symptomatic of some misspecification in the formulation of the targeting problem or is it a cruel but acceptable feature of benign policy-making?

2. Modelling the targeting problem

The population, of size unity, is divided into two mutually exclusive and exhaustive groups, A and B, with a proportion θ in group A. The elements of these groups are referred to as individuals, though they can equally well be thought of as households: group A might consist of married couples, for instance, and group B of single people. Individuals cannot affect the group to which they belong. Pre-transfer income, y , is exogenous (precluding labour supply effects) and distributed continuously with support in $[y, \bar{y}]$.

⁴ An earlier version of this paper (with the same title, and available from the author on request) also considered the possibility of another apparent perversity: might it be optimal to over-compensate a group for an increase in its needs? Roughly speaking, it turns out that whilst this is indeed possible when the headcount measure is used, fairly mild conditions suffice to rule it out more generally.

The information available to the government is assumed throughout to include the group to which each individual belongs, the population weight θ and the group-specific income distributions $F_i(y)$ ($i=A,B$); in Section 4 this information set is expanded to include the income of each individual. The distribution functions F_i are taken to be twice continuously differentiable, the corresponding density functions being denoted $f_i(y)$. The government arranges transfer/taxes that result in net income of $n_i(y)$ for a member of group i with pre-transfer income y .

The object of policy is to minimise an aggregate poverty index P that is a weighted sum of sub-indices P_i measuring poverty within the two groups:

$$P = \lambda P_A + (1-\lambda)P_B, \quad \lambda \in (0,1). \quad (2.1)$$

The weight λ is taken as parametric. It may coincide with the group's population share, reflecting additive decomposability of the aggregate index; but other interpretations are also possible. The important feature is that λ does not vary with either income patterns or the poverty lines introduced in a moment (and even this is essentially a matter of simplicity). Assuming net income to be non-decreasing in pre-transfer income, within-group poverty is measured relative to a group-specific poverty line z_i by an index⁵ of the form

$$P_i = \int_{z_i}^{\psi_i} p[n_i(y), z_i] f_i(y) dy, \quad (2.2)$$

where $\psi_i = \max\{y | n_i(y) \leq z_i\}$ is the largest pre-transfer income at which an individual in group i remains post-transfer poor and $p[n, z]$ measures the extent of an individual's deprivation, with $p[n, z] \geq 0$ and $p[n, z] = 0$ whenever $n > z$. We refer to $p[.]$ as the *deprivation function*. It is assumed that the same deprivation function is applied to both groups; at the expense of an additional subscript, the results below generalise trivially. The deprivation function is taken throughout to be twice continuously differentiable (except perhaps at $n = z$) and, denoting derivatives⁶ by subscripts, to satisfy:

⁵To ensure that the integral in (2.2) exists, it is assumed throughout that $n_i(y) > 0$, $i=A,B$.

⁶Those with respect to n at $n=z$ are left-hand derivatives.

$$(a) \quad p_n[n, z] \leq 0 \quad (2.3a)$$

$$(b) \quad p_z[n, z] \geq 0 \quad (2.3b)$$

$$(c) \quad p_{nn}[n, z] \geq 0 \quad (2.3c)$$

$$(d) \quad p[z, z] = \kappa \geq 0, \text{ independent of } z. \quad (2.3d)$$

Properties (a) and (b) simply mean that deprivation is never increased by an increase in net income or by a reduction in the poverty line. The convexity property (c) corresponds to Sen's (1979) transfer principle, ensuring that poverty is never reduced by a small transfer of income from one poor individual to another but richer one (except, perhaps, if it takes the latter over the poverty line). Property (d) is less familiar, but seems uncontentious: the usual poverty measures count deprivation at the poverty line as either zero or one.

The additively separable form of the aggregate index (2.1)-(2.2) rules out a number of indices that have appeared in the literature, such as those which, following Sen (1976), incorporate Gini-type components. Nevertheless, the formulation captures a wide class of measures, including those analysed in Atkinson (1987). Note too that it satisfies the sub-group consistency axiom of Foster and Shorrocks (1989).

It will be instructive at various points to put more structure on the poverty measure. In particular, it will prove important to distinguish between *relative* and *absolute* measures of deprivation:⁷ the former depend only on the ratio of net income to the poverty line, so that $p[n, z]$ can be written in the form $\pi(n/z)$; the latter take $p[n, z]$ to be of the form $\pi(n-z)$. In each of these cases, the conditions of (2.3) are met if π is non-increasing and convex. Foster and Shorrocks (1989) show that any poverty index satisfying (2.2) and (2.3) that belongs to both of these

⁷ This distinction is closely related to that between relative and absolute poverty indices drawn by Blackorby and Donaldson (1980). Both, it should be emphasised, are quite distinct from the traditional dichotomy between absolute and relative concepts of poverty, which - put crudely - revolves around the substantive issue of whether or not poverty lines should respond to changes in the general level of material well-being. Adopting an absolute concept of poverty, for instance, one might believe that z should not be changed if all individuals' real net incomes were to double; but one might nevertheless choose to measure the poverty that remains using a relative deprivation measure.

classes must be an increasing transform of the headcount ratio, for which

$$p[n, z] = 1 \quad \text{when } n \leq z . \quad (2.4)$$

Despite widespread criticism, this remains a popular measure in both empirical work and policy discussion. The same is true of the per capita income gap ratio measure, for which

$$p[n, z] = [z-n]/z \quad \text{when } n \leq z . \quad (2.5)$$

Another important special case is the measure of Foster, Greer and Thorbecke (1984, henceforth FGT), which has dominated the recent targeting literature. This measures deprivation as

$$p[n, z] = [(z-n)/z]^\alpha \quad \text{when } n \leq z \quad (2.6)$$

where $\alpha \geq 1$ (for convexity). Reference will also be made to the index of Watts (1968), for which

$$p[n, z] = \ln(z/n) \quad \text{when } n \leq z \quad (2.7)$$

and to that of Clark, Hemming and Ulph (1981, henceforth CHU) for which

$$p[n, z] = [1 - (n/z)^{\beta+1}]/(\beta+1) \quad \text{when } n \leq z \quad (2.8)$$

where $\beta \in (-\infty, 0) - \{-1\}$; note that the former corresponds to $\beta \rightarrow -1$ in the latter.

The government is constrained in its attempt to minimise poverty by the requirement that the cost of the transfer scheme not exceed some fixed amount R , the sign of which is unrestricted. This revenue constraint is assumed to bind at the optimum, and so is imposed as the equality

$$R = \theta \int_{\underline{y}}^{\bar{y}} \{n_A(y) - y\} f_A(y) dy + (1-\theta) \int_{\underline{y}}^{\bar{y}} \{n_B(y) - y\} f_B(y) dy . \quad (2.9)$$

No attempt will be made to characterise the fully optimal net income schedules. Rather the approach is to impose *a priori* restrictions on the schedules that can be chosen - restrictions that mimic widely advocated strategies for fundamental reform - and then to analyse the role of poverty lines in designing conditionally optimal patterns of targeting.

3. Universality

We begin with the case in which the only instruments available to the government are poll subsidies S_i ($i=A,B$), constant within but possibly different across the two groups. Thus the net income schedules are in this section restricted to the form

$$n_i(y) = y + S_i, \quad i=A,B. \quad (3.1)$$

This would be the only feasible deterministic strategy if the governments' information set were no richer than described in the preceding section. This in turn may be a reasonable approximation to the position in many developing countries, where both administrative problems and the importance of non-market transactions combine to make the accurate measurement of original income problematic. More generally, such a structure of poverty relief captures in its purest form the universality principle associated with Beveridge: benefit receipt is contingent on an individual's type but independent of their income.

The example given in the Introduction shows that the paradox of targeting can indeed occur when the government adopts a strategy of this kind:⁸ the optimal response to an increase in the needs of a particular group may be to reduce its poll subsidy. The purpose now is to investigate the conditions under which it can arise more generally and precisely.

Substituting (3.1) into (2.1)-(2.2) and using the revenue constraint to eliminate⁹ S_B gives aggregate poverty as a function of the benefit paid to and (for the comparative statics to follow) the poverty line of group A. In obvious notation (and recalling that $\kappa = p[z,z]$), the first order condition for the targeting problem under universality is then that

⁸Strictly, the example in the Introduction assumes a discrete distribution of income rather than, as here, a continuous one; but it is straightforward to construct similar examples for the continuous case.

⁹Using (3.1) in (2.9) gives $S_B = \bar{R} - \bar{\theta}S_A$, where $\bar{R} = R/(1-\theta)$ and $\bar{\theta} = \theta/(1-\theta)$.

$$P_S(S_A^*, z_A) = \lambda \left(\int_{\underline{y}}^{\psi_A} p_n[y+S_A^*, z_A] f_A(y) dy - \kappa f_A(z_A - S_A^*) \right) - (1-\lambda) \theta \left(\int_{\underline{y}}^{\psi_B} p_n[y+S_B^*, z_B] f_B(y) dy - \kappa f_B(z_B - S_B^*) \right) = 0. \quad (3.2)$$

This simply equates across the two groups the amount by which poverty in each would be reduced by a small increase in its poll subsidy (adjusted for revenue cost and weight in the aggregate index), this reduction comprising both an intra-marginal effect on the deprivation of those who remain poor and an effect on the extensive margin through those taken out of poverty altogether.

When then will optimal targeting require that an increase in the needs of group A be met by a cut in the poll subsidy it receives? Applying the implicit function theorem to (3.2) and assuming the second order condition to be satisfied,¹⁰ dS_A^*/dz_A has the same sign as

$$\Delta = \kappa f'_A(z_A - S_A^*) - \int_{\underline{y}}^{\psi_A} p_{nz}[y+S_A^*, z_A] f_A(y) dy - p_n[z_A, z_A] f_A(z_A - S_A^*), \quad (3.3)$$

(the prime indicating the derivative of the density). The final term in (3.3) is non-negative by (2.3a), and so pushes towards a positive relationship between needs and benefits; the signs of the other two terms, however, are ambiguous. Clearly the possibility of a paradox depends in a complex way on the shapes of both the deprivation function and the distribution of income. We consider the role of each in turn.

Suppose first that the deprivation function is of the absolute form, so that $p[n, z] = \pi(n-z)$. Then

¹⁰ The second order condition is that $\lambda \Omega_A + (1-\lambda) \bar{\theta}^2 \Omega_B$ be strictly positive at S_A^* , where

$$\Omega_i = \int_{\underline{y}}^{\psi_i} p_{nn}[y+S_i, z_i] f_i(y) dy - p_n[z_i, z_i] f_i(z_i - S_i) + \kappa f'_i(z_i - S_i).$$

Except, when $\kappa > 0$, it is sufficient for this that (2.3a) hold with strict inequality over some range of observed post-transfer incomes. For practical purposes it is thus only in the case of the headcount ratio - which will be dealt with separately below (see footnote 11) - that the second order condition is potentially problematic.

$$-p_{nz}[n, z] = p_{nn}[n, z] \quad (3.4)$$

is non-negative by the convexity condition (2.3c), and it is immediate from (3.3) that:

PROPOSITION 1: *The paradox cannot arise under universality if the deprivation function is of the absolute form and $\kappa = 0$.*

This is easily explained. With deprivation measured in absolute terms a unit increase in the poverty line of group A is precisely equivalent to a unit reduction in all its members' net incomes. By convexity, this increases (strictly, does not reduce) the marginal deprivation $-p_n[n, z]$ of the poor in A and so strengthens (does not weaken) the intramarginal gains from re-targeting resources towards them; and with $\kappa=0$ there is no effect at the extensive margin to offset this.

Matters are very different when poverty is measured in relative terms. In this case $p_n[n, z] = \pi'(n/z)/z$, so that an increase in the poverty line affects marginal deprivation in two ways. One is equivalent to a uniform proportionate reduction in net incomes; this, by convexity, increases the gain from transferring 'equivalised' net income n/z to those in the group concerned. The second effect, however, is to increase the revenue cost of supporting equivalised incomes (doubling z , for instance, doubles the cost of raising n/z); and this points towards a retargeting away from the group whose needs have increased. This suggests that an increase in z_A might lead to a reduction in S_A^* if the first of these effects is sufficiently weak. More precisely, since in this case

$$p_{nz}[n, z] = -\{\pi''(n/z)(n/z) + \pi'(n/z)\}/z^2, \quad (3.5)$$

with $\pi'' \leq 0$ by (2.3c), (3.3) gives:

PROPOSITION 2: *If the deprivation function is of the relative form with $\kappa = 0$, the paradox can arise under universality only if*

$$-\frac{d\ln\{-\pi(x)\}}{d\ln(x)} = -\frac{\delta\ln(-p_n)}{\delta\ln(n)} < 1 \quad \text{for some } x = n/z < 1. \quad (3.6)$$

The paradox thus requires marginal deprivation to be inelastic in net

income: only then can the cost argument outweigh the increased effectiveness of equivalised transfers. Put another way, the paradox can be ruled out if the poverty index is sufficiently averse to inequality amongst the poor. Interestingly, the critical elasticity in (3.6) corresponds precisely to the parameter $-\beta$ in the CHU index (2.8), so that the paradox would then immediately be precluded by taking $\beta < -1$. Nor can the paradox arise when using the Watts index, since the critical elasticity is then exactly unity.

Turning to the second of the broad factors influencing the optimal pattern of targeting - the distribution of pre-transfer income - integrating by parts in (3.3) and using the implication of (2.3d) that $p_n[z, z] = -p_z[z, z]$ gives

$$\Delta = \kappa f'_A(z_A - S_A^*) + \int_{\underline{y}}^{\psi_A} p_z[y + S_A^*, z_A] f'_A(y) dy + p_z[y + S_A^*, z_A] f_A(y) . \quad (3.7)$$

Recalling (2.3b), it is then immediate that:

PROPOSITION 3: *Whatever the form of the deprivation function, the paradox cannot occur under universality if the density of post-transfer income is increasing everywhere below the poverty line.*

This in turn gives:

COROLLARY 3.1: *Whenever the (within-group) distribution of income is unimodal, the paradox can occur under universality only if, at the initial optimum, modal net income is below the poverty line.*

To see how the paradox might arise when the mode lies below the poverty line, note first that an increase in z_A then reduces the number of individuals exactly at the poverty line and so, when $\kappa > 0$, operates on the extensive margin in such a way as to weaken the case for targeting towards A. This effect is clearest in the case of the headcount index. Using (2.4) and its implications in (3.3), and assuming the second order condition to be satisfied,¹¹ the paradox then arises iff $f'_A(z_A - S_A^*) < 0$: the modal condition

¹¹ This requires that $f'_A + (1-\lambda)(\theta)^2 f'_B > 0$. Note that the paradox cannot then

of Corollary 3.1 is in this case sufficient as well as necessary. More generally, an increase in the poverty line also affects the marginal deprivation of those who remain poor, corresponding to the term $\int p_n^* f dy$ in (3.2). A similar but more complex effect is at work here too. Assume for brevity that $\kappa = 0$. Then (3.3) implies that the paradox can only arise if $p_{nz} > 0$ over some interval. When this is so (and assuming strict convexity of the deprivation function in n), an increase in the poverty line reduces the level of pre-transfer income associated with the corresponding levels of marginal deprivation. Imagine that this condition holds over some range immediately below the poverty line, and that the density is falling there. Then one effect of increasing z_A will be to increase the number of individuals suffering 'small' levels of marginal deprivation. It may be that it also increases the number experiencing 'large' marginal deprivation further down the distribution. But if the first of these effects dominates - and this will depend on the form of the deprivation function - then the total intra-marginal poverty reduction achieved by a small payment to all those in A will fall, making it optimal to retarget towards the group whose poverty line is unchanged.

The broad implication of Corollary 3.1 is that the paradox is unlikely to occur under universality unless poverty is pervasive amongst the group concerned. The prevalence of poverty will in turn depend on, amongst other things, the precision with which groups at particular risk of poverty can be singled out for differential treatment, so that this result enables no simple generalisation on the likely practical significance of the paradox. It is worth noting, however, that it is not hard - particularly in less developed countries - to find policy contexts in which the modal net incomes of some groups lie below conventional poverty lines: Ravallion (1988), for instance, cites data for three Indian villages in which the estimated modal income falls short of the poverty line. This suggests that the possibility of paradoxes when using standard poverty measures cannot be dismissed lightly, though it should be emphasised that the modal condition of Corollary 3.1 is not sufficient, and, moreover, is required to hold after

arise for both groups, a property that can be shown to be common to all absolute poverty indices.

optimal targeting of universal benefits.

As with the CHU and headcount indices already discussed, particular forms of poverty index generate convenient necessary and sufficient conditions for the occurrence of the paradox. Of particular interest is the widely-used FGT index. For $\alpha > 1$, this is readily shown to yield the paradox iff

$$\alpha P_A^{\alpha-1} > (\alpha-1) P_A^{\alpha-2}, \quad (3.8)$$

where P_A^α denotes the FGT index for A evaluated at α . The issue thus turns on a comparison of FGT indices evaluated at values of the 'deprivation aversion' parameter lower than that used in the minimand: the calculation of sub-group indices would itself reveal the scope for paradoxes. For the common choice $\alpha = 2$ (measured poverty then reflecting, inter alia, the coefficient of variation of income amongst the poor), it is straightforward to show that (3.8) takes a particularly simple form: the paradox arises iff, at the optimum, the mean net income of the poor is less than half the poverty line.

4. Means-testing

Alternative strategies for poverty alleviation become available if, as is now assumed, the government can costlessly observe pre-transfer incomes. This section examines the relationship between needs and benefits under one such strategy, in which the government simply provides guaranteed minimum incomes - possibly at different rates - to the two groups. The benefit system thus acts purely as a safety net, the role originally envisaged by Beveridge for National Assistance (now Income Support) in the UK. Such a scheme corresponds to an extreme form of means-testing, since benefit recipients face a marginal tax rate of 100 per cent. This structure is unlikely to be fully optimal even in the present simple model.¹²

¹² It may indeed even be inferior to universality: using the headcount measure, for instance, universality will dominate when some of the poor are close to the poverty line and many others far below it. For poverty indices that are sufficiently averse to deprivation amongst the poor, one might expect means-testing to be preferred. More complex forms of income

Nevertheless, it serves to characterise a strategy that continues to play a central role both in practical social security design, not least in the UK, and in discussion of its reform.

If sufficient resources were available, the optimal policy of this kind - referred to simply as means-testing - would be to guarantee incomes at (or, if $\kappa > 0$, slightly above) the poverty lines z_i . More generally, it becomes necessary to distinguish between the poverty lines that the government uses to evaluate the extent of poverty, z_i , and those that it sets as guaranteed incomes, c_i . Thus net incomes are now:

$$n_i(y) = \min\{y, c_i\} . \quad (4.1)$$

Endogenising guaranteed income levels in this way has some appeal in terms of political economy: politicians in the UK, for instance, have been alive to the fact that raising levels of income support tends to increase measured poverty and anxious (at least when in office) to argue that this in itself represents no real change in the extent of the underlying policy problem.

Under means-testing, group-specific poverty becomes

$$P_i = p[c_i, z_i]F_i(c_i) + \int_{c_i}^{z_i} p[y, z_i]f_i(y)dy, \quad i=A, B \quad (4.2)$$

and the targeting problem that of choosing the official poverty lines c_A and c_B to minimise the aggregate poverty index defined by (2.1)-(2.2) and (4.2) subject to the revenue constraint (2.9). In this context the paradox of targeting would correspond to a negative relationship between the optimal official poverty line for, say, group A, denoted c_A^* , and the true poverty line z_A . The example given in the Introduction again suffices to show that this can indeed happen: this follows on noting from (4.2) that when the headcount ratio is used it can never be optimal to set an official poverty line below the true.

relation, however, will generally target even more effectively. But little would be gained by pursuing the question of full optimality in the present framework, ignoring as it does both incentive effects (analysed in the context of poverty minimisation in Kanbur, Keen and Tuomala (1990)) and considerations of stigma and hassle attached to benefit receipt (Cowell (1986)).

For the more general case, it is routine to show that when $F_B(c_B^*) > 0$, so that some of those in group B receive benefit at the optimum, the minimisation of aggregate poverty requires that¹³

$$P_c(c_A^*, z_A) = \{\lambda p_n[c_A^*, z_A] - \theta(1-\theta)^{-1}(1-\lambda)p_n[c_B^*, z_B]\}F_A(c_A^*) = 0 \quad (4.3)$$

The intuition here is that raising the official poverty line of group A by £1 increases by £1 the net incomes of the $\theta F_A(c_A)$ individuals in the group receiving benefit, which reduces the deprivation of each by $p_n[c_A, z_A]$; condition (4.3) simply equates the impact of this on aggregate poverty with that of an equal-cost increase in c_B .

It is immediate from (4.3) that the paradox is also inescapable whenever the income gap ratio measure (2.5) is used. Assuming, for brevity, that $\theta = \lambda$ (so that the weights in the aggregate poverty index are simply population weights), all available resources will be allocated towards group A (at least until its poverty has been eradicated) iff $z_A < z_B$. The reason is clear. Under means-testing it is possible to transfer income directly between those on the official poverty lines of the two groups. Differences in group size and weighting apart, the aggregate equivalised poverty gap will then always be reduced by retargeting towards the group in which the price of increasing equivalised income, z_i , is lower. Increasing z_A from a sufficiently low level to a sufficiently high one will thus always cause a discontinuous retargeting towards group B.

The headcount and income gap ratio measures have the common feature of being insensitive to the distribution of income amongst the poor. When the deprivation function is strictly convex (which ensures that the second order condition for the targeting problem is satisfied), routine application of the implicit function theorem to (4.3) gives a strikingly simple necessary and sufficient condition for the occurrence of the paradox:

PROPOSITION 4: *Assuming some members of both groups to be post-transfer poor at the initial optimum and the deprivation function to be strictly*

¹³ Equation (4.3) follows on differentiating (4.2), noting that the revenue constraint in this case implies $dc_B/dc_A = -\bar{\theta}F_A(c_A)/F(c_B)$.

convex, the paradox arises under means-testing ($dc_A^*/dz_A < 0$) if and only if

$$p_{nz}[c_A^*, z_A] > 0. \quad (4.4)$$

The explanation is straightforward. When (4.4) holds, an increase in the needs of group A reduces the impact on its deprivation of a marginal increase in the official poverty line, weakening the case for targeting towards it. Note that in this case - in sharp contrast to the corresponding condition ($\Delta < 0$) under universality - the distribution of pre-transfer income has no direct bearing on the possibility of the paradox (though it will of course have an indirect effect, influencing official poverty lines through the revenue constraint).

Recalling (3.4), it is immediate from Proposition 4 that:

PROPOSITION 5: *The paradox cannot arise under means-testing if the deprivation function is of the absolute form and strictly convex.*

The reasoning behind this is similar to that given for Proposition 1: when deprivation is measured in absolute terms, a £1 increase in needs is equivalent to a £1 reduction in the official poverty line and so, by the second order condition, calls for an increase in the latter.

From (3.5) we also have:

PROPOSITION 6: *When the deprivation function is of the relative form and strictly convex, the paradox arises under means-testing if and only if*

$$-\frac{\delta \ln(-p_n)}{\delta \ln(n)} < 1 \quad \text{at } n = c_A^*. \quad (4.5)$$

This elasticity is precisely as in Proposition 2 above for universality; here, however, the restriction is required to hold only at the official poverty line and is sufficient as well as necessary.

The condition in (4.5) will often be easy to check. When the CHU index is used, for example, Proposition 6 implies that the paradox will arise iff

$\beta \in (-1, 0)$. Watts' index, interestingly, emerges as a borderline case in which official poverty lines are optimally independent (locally around an interior solution) of true needs. Using an FGT index (with $\alpha > 1$ for strict convexity), the paradox occurs iff $\alpha < z/c^*$; and so is more likely to arise, loosely speaking, the lower is aversion to inequality amongst the poor and the deeper is poverty at the initial optimum. At least for the UK, the orders of magnitude usually encountered in empirical work suggest that the paradox will often be of limited practical importance in developed countries when the FGT index is used. Morris and Preston (1986), for example, consider the extent of poverty when true poverty lines are up to 120 per cent of the corresponding Supplementary Benefit levels: the paradox could then be ruled out if the latter were the product of an optimally targeted means-testing programme with α above 1.2, a value rather lower than is commonly employed. But a firmer feel for the policy significance of the paradox again requires that one first go through the optimal targeting exercise.

5. Summary and conclusions

The precise conditions under which optimally-targeted benefits are inversely related to needs depend on the nature of the poverty alleviation strategy being pursued. Nevertheless, the analysis points to three broad factors which, loosely speaking, make the occurrence of the paradox more likely. The first is that the extent of an individual's deprivation be measured in terms of the ratio of their net income to the poverty line (rather than, in particular, in terms of the absolute difference between the two): £1 then buys a greater reduction in poverty if targeted to a less needy individual. The second is that the rapidity with which an individual's measured deprivation increases as their net income falls not be too sensitive to the initial level of that income: for an increase in needs is akin to a reduction in net income, and so in these circumstances will do relatively little to enhance the effectiveness of transfers in reducing deprivation. The third is that poverty be pervasive within the group affected.

Of wider importance than the details of these results, however, is the

general issue that they raise: How should one respond to the paradox? Its practical significance remains unclear: detailed simulation studies are needed to give a firm impression of whether or not the paradox is likely to arise with commonly-used micro data, common specifications of poverty lines and common choices of poverty index. But the essential issue that has to be faced is conceptual rather than empirical: without further restrictions on the structure of the targeting problem than are conventionally imposed, the paradox simply cannot be precluded. Indeed there exist orthodox poverty indices (of the CHU type, with $\beta \in (-1, 0)$) and poverty alleviation strategies (means-testing of the kind above) such that it is a certainty. How one responds to the paradox must thus ultimately be a matter of methodology.

There are three main possibilities. First, the paradox might be regarded as symptomatic of a deeper malaise in the specification of the policy problem. Stern (1987), for instance, argues that the critical role of the poverty line in defining poverty measures means that - whatever their descriptive value - they are singularly inappropriate as minimands for policy design. In the present context, however, the culprit - if there is one - is less the presumed objective function than it is the parameterisation of needs through simple equivalence scales: this is readily shown to create scope for similar possibilities when pursuing standard welfarist objectives. Subtler characterisations of differences of needs are indeed possible, along the lines of Atkinson and Bourguignon (1987), but in themselves seem unlikely to remove the fundamental source of the paradox: £1 may do most good in alleviating deprivation if given to those whose needs are relatively low.¹⁴ Second, one might accept the broad policy framework but resist the implication. The results derived above might thus be used to argue for further acceptability restrictions on poverty indices. In all the circumstances considered here, for example, the paradox can be ruled out by requiring that the measure of deprivation be of the absolute form, strictly convex and treat as non-deprived those exactly on the poverty line. Those circumstances are limited, however, to two particular forms of poverty alleviation strategy. The question then is whether this or any other set of

¹⁴ This would indeed seem to be suggested, for instance, by the recent work of Lambert (1990).

conditions on the poverty index suffices to preclude the paradox when the form of the poverty alleviation strategy is unrestricted. The third response is both the hardest and the most facile. It is to take the essence of the paradox at face value: with value judgements left to the policy-makers, efficient targeting may be cruel to the very needy.

References

- Atkinson, A.B. (1987). 'On the measurement of poverty.' Econometrica, 55, 749-764.
- Atkinson, A.B. and Bourguignon, F. (1987). 'Income distribution and differences in needs.' In Feiwel, G.R. (ed) Arrow and the Foundations of the theory of Economic Policy. New York: MacMillan.
- Beath, J., Lewis, G. and Ulph, D. (1988). 'Policy targeting in a new welfare framework with poverty.' In Hare, P.G. (ed) Surveys in Public Sector Economics. Oxford: Basil Blackwell.
- Besley, T. (1990). 'Means-testing versus universal provision in poverty alleviation programmes.' Economica, 57, 119-129.
- Besley, T. and Kanbur, R. (1988). 'Food subsidies and poverty alleviation.' Economic Journal, 98, 701-719.
- Blackorby, C. and Donaldson, D. (1980). 'Ethical indices for the measurement of poverty.' Econometrica, 48, 1053-1060.
- Clark, S., Hemming, R. and Ulph, D. (1981). 'On indices for the measurement of poverty.' Economic Journal, 91, 515-526.
- Cowell, F.A. (1986). 'Welfare benefits and the economics of take up.' ESRC Programme on Taxation, Incentives and the Distribution of Income Discussion Paper No.89, London School of Economics.
- Foster, J., Greer, J. and Thorbecke, E. (1984). 'A class of decomposable poverty measures.' Econometrica, 52, 761-766.
- Foster, J. and Shorrocks, A.F. (1989). 'Subgroup consistent poverty indices.' Mimeo, University of Essex.
- Glewwe, P. (1988). 'Household equivalence scales and the measurement of inequality: Some paradoxical implications.' Mimeo, World Bank.
- Kanbur, R. (1987). 'Transfers, targeting and poverty.' Economic Policy, 4, 112-136 and 141-147.
- Kanbur, R. and Keen, M.J. (1989). 'Poverty, incentives and linear income taxation.' In Dilnot, A.W. and Walker, I. (eds) The Economics of Social Security. Oxford: Oxford University Press.
- Kanbur, R., Keen, M.J. and Tuomala, M. (1990). 'Optimal non-linear income taxation for the alleviation of income poverty.' Mimeo, Queens University.
- Lambert, P.J. (1990). 'Redistribution through the income tax, or: Robin Hood and the groves of the forest.' Institute for Fiscal Studies Discussion Paper W90/13

- Morris, N. and Preston, I. (1986). 'Inequality, poverty and the distribution of income.' Bulletin of Economic Research, 38, 277-344.
- Ravallion, M. (1988). 'Expected poverty under risk-induced welfare vulnerability.' Economic Journal, 98, 1171-1182.
- Ravallion, M. and Chao, K. (1989). 'Targeted policies for poverty alleviation under imperfect information: Algorithms and applications.' Journal of Policy Modeling, 11, 213-224.
- Sen, A.K. (1976). 'Poverty: An ordinal approach to measurement.' Econometrica, 44, 219-231.
- Sen, A.K. (1979). 'Issues in the measurement of poverty.' Scandinavian Journal of Economics, 81, 285-307.
- Stern, N.H. (1987). 'Discussion.' Economic Policy, 4, 136-140.
- Watts, H.W. (1969). 'An economic definition of poverty.' In Moynihan, D.P. (ed) On Understanding Poverty. New York: Basic Books.