AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

OPTIMAL NON-LINEAR INCOME TAXATION FOR THE ALLEVIATION OF INCOME
POVERTY

By

Ravi Kanbur
University of Warwick and the World Bank

Michael Keen
University of Essex and Queens University

and

Matti Tuomala
University of Jyväskylä

No. 368

November 1990

Abstract: This paper is concerned with the optimal use of income information in the design of tax/transfer systems to alleviate poverty. The issue is one of optimal non-linear income taxation, but using a non-welfarist objective function that seems to accord well with the common concerns of policy debate: an income-based poverty index. We show that one of the key results of the welfarist literature is overturned: if it is desirable for everybody to work, the optimal marginal tax rate on the very poorest individuals is strictly negative. More generally, it is argued that the non-welfarist perspective points towards lower marginal tax rates in the lower part of the distribution than does the welfarist. Numerical simulations suggest, however, that this effect is of limited quantitative significance. Using conventional functional forms and parameter values, optimal marginal tax rates on the poor are in the 60-70% range.

Acknowledgements: We are grateful to Tim Besley for many very helpful discussions on these and related issues. Our collaboration began at the 1988 Warwick Summer Workshop, and we thank the Department of Economics at Warwick for its hospitality. Views and errors are ours.

This paper is circulated for discussion purposes only and its contents should be considered preliminary

1. Introduction

In both developed and developing countries there has recently been much discussion of 'targeting' in the design of social security and income transfers. By this is meant the structuring of tax and transfer programmes so as to ensure that resources are concentrated on the poor with leakages to the non-poor minimised. Two dimensions of the targeting problem are of paramount interest. One is the issue of how best to use non-income information on various contingencies - such as labour market status (unemployment, retirement), demographic characteristics or region of residence - to differentiate transfers. Aspects of this have been analysed by Besley and Kanbur (1988) in a developing country context and by Kanbur and Keen (1989) in a model of group-specific linear tax schedules. This work has typically assumed income-relation to be infeasible or restricted it to take a very simple form, so precluding any consideration of the second general issue. This is the appropriate nature and degree of income-relation ('means testing') in determining the extent of an individual or household's transfer/tax. It is to this issue that the present paper is addressed.

In the absence of incentive effects the design of a perfectly-targeted non-linear income tax is a trivial exercise. Having established a poverty line, simply give everyone who is initially below it exactly that transfer needed to bring them just above it.¹ Such a scheme involves no leakages. But once incentive effects are admitted an obvious and familiar difficulty arises. Since the scheme implies a marginal tax rate on those below the poverty line of 100%, the poor have no incentive to earn income. Their rational labour supply decisions would then be liable to greatly increase the revenue cost of alleviating their poverty. Incentive effects may thus force one to rule out marginal rates on the poor of 100%. The questions of precisely how high or low those rates should be, and of how they should vary with income, then become considerably more complex. Opinion seems to vary widely. On the one hand, many social security schemes observed in practice face the poor with extremely high effective marginal tax rates (inclusive of

¹ This assumes that the resources available are sufficient to entirely eradicate poverty; a distinct set of issues arises if they are not.

the taxation implicit in benefit withdrawal). In the UK, for instance, while the Fowler reforms eliminated the poverty trap in the extreme form of marginal rates in excess of 100% they are likely to have had the effect of extending the range of incomes over which the poor face very high marginal rates (Dilnot and Stark (1990)). In the US, the AFDC programme alone has implied marginal rates in the order of 70%. Some reform proposals, in contrast - such as those seeking to restore a reduced rate band for income tax of the kind implemented in the UK from 1978 to 1980 - call for relatively low marginal tax rates at the bottom of the distribution, presumably in the hope of encouraging the low-paid to help themselves by their own efforts. Such fundamentally different strategies for poverty alleviation raise clearly the underlying issue of principle: What is the shape of the tax-transfer schedule that does most to reduce poverty?

It seems clear that this shape is unlikely to be a simple one. In particular, restricting attention to linear tax schedules would instantly rule out of court much of the interesting area of debate. Thus while in Kanbur and Keen (1989) the analysis is restricted to linear schedules - in order to concentrate on the use of non-income information - here we take the opposite approach of assuming individuals to be identical in all but their pre-tax wage rates (so removing any role for non-income information) whilst allowing for fully non-linear taxation.

The question being asked here is of course closely related to that addressed in the literature, initiated by Mirrlees (1971), on optimal non-linear income taxation. In this framework, the issue of work incentives is tackled directly by modeling individuals as choosing between work and leisure given the tax-transfer schedule they face. The government then chooses a schedule which maximises a social welfare function based on individuals' welfare, that is on the utility they derive from their consumption-leisure bundles. There is, however, a striking and fundamental dissonance between this 'welfarist' literature and the tone of the policy debates around, for instance, the Fowler reforms. For while the welfarist literature takes into account the values of both net income and leisure in the individual welfare functions, the policy discussion focusses almost exclusively on incomes. It is the consequences of reform for the incomes of the poor - the money in their

pockets, not something akin to money metric measures of their welfare - that is commonly debated and analysed. Their valuations of the associated consumption-leisure bundle are typically ignored. Even when work incentives are discussed explicitly it is the implications for government revenue and individual incomes that are paramount; little weight is typically given to such disutility as the poor experience from working. We do not attempt to explain why this is so, though that is an interesting and important question in itself. Rather our purpose is to take the concerns of the policy debate as given and examine their implications: we seek to investigate, both qualitatively and quantitatively, the central features of the non-linear tax schedules that minimise income-based measures of poverty.

The present analysis is thus squarely in the 'non-welfarist' camp. This is not to say that the welfarist viewpoint is misplaced. It is simply to suggest that an alternative which appears more closely related to the normal terms of practical policy formulation and evaluation merits close analysis. Despite growing interest in non-welfarist approaches to policy problems, there has been little discussion of the robustness of welfarist results to the adoption of non-welfarist perspectives (an exception being Seade (1980)). Thus while our central aim is to address the policy debate on income-testing our results can also be seen as contributing to the evaluation of the sensitivity of qualitative results of the traditional literature to variations in one its fundamental premisses.

The plan of the paper is as follows. Section 2 reviews the key qualitative features of the welfarist literature on optimal non-linear income taxation. Section 3 contains our main analytical results, on the pattern of marginal rates that minimises some measure (within a wide class) of poverty defined on net incomes (or, equivalently, consumption). This pattern turns out to be qualitatively very different from that which emerges from the traditional welfarist approach. These distinctive results pertain, however, primarily to a neighbourhood at the lower end of the schedule, so that it is important to investigate the quantitative pattern of the schedule over the whole range. For this numerical simulation is needed. Section 4 presents and discusses such results. Section 5 concludes, and an Appendix outlines both the derivation of the main analytical results and the method of the simulations.

2. Lessons from the welfarist literature

We assume there to be a continuum of consumers, each having preferences $u(x,y)$ defined over consumption x and hours worked y , with $u_x > 0$ and $u_y < 0$ (subscripts indicating partial derivatives). Individuals differ only in the pre-tax wage n they can earn, which is distributed across the population with continuous density $f(n)$ on support $[\underline{n}, \bar{n}]$. Writing gross income as $z = ny$ and defining

$$s(x,z,n) = -u_y(x,z/n)/nu_x(x,z/n) > 0, \quad (2.1)$$

preferences are taken to satisfy the further restriction that

$$s_n < 0. \quad (2.2)$$

This is Assumption B of Mirrlees (1971) and the Agent Monotonicity assumption of Seade (1982). It implies that indifference curves in consumption-gross income space become flatter the higher is an individual's wage rate, which in turn ensures that both consumption and gross earnings increase with the wage rate. Each individual maximises utility by choice of hours worked, solving

$$\max_{x,y} u(x,y) \quad \text{subject to} \quad x = ny - T(ny). \quad (2.3)$$

This gives rise to consumption, gross income and maximised utility denoted $x(n)$, $z(n)$ and $v(n)$ respectively.

The problem of a welfarist government is then to

$$\max_T \int_{\underline{n}}^{\bar{n}} W[v(n)]f(n)dn \quad \text{subject to} \quad \int_{\underline{n}}^{\bar{n}} T[z(n)]f(n)dn = R \quad (2.4)$$

where $W[.]$, which is taken to be concave and increasing, gives the social valuation of utility and R denotes the revenue requirement. The self-selection constraint imposed by individuals' optimisation, implicit in (2.4), is conveniently characterised in terms of the necessary condition for (2.3), which can be written as

$$u_x[1-t(z)] + u_y/n = 0, \quad (2.5)$$

where $t(z) = T'(z)$ denotes the marginal tax rate at z .

Omitting details - which are analogous to those given in Appendix A for the analysis of poverty minimisation in the next section - the first order conditions for the welfarist problem imply a pattern of marginal rates

satisfying, for all n such that labour supply is strictly positive,²

$$t[z(n)] = - \frac{\mu(n)u_x s_n}{\lambda f(n)} \quad (2.6)$$

where λ is the multiplier on the revenue constraint in (2.4) and (denoting du/dn by u_n)

$$\mu(n) = \int_{\underline{n}}^n (W' u_x - \lambda) (1/u_x) \exp\left(-\int_p^n (u_{nx}/u_x) dm\right) f(p) dp \quad (2.7)$$

is that on the incentive compatibility constraint. This latter satisfies the transversality conditions

$$\mu(\underline{n}) = \mu(\bar{n}) = 0. \quad (2.8)$$

Using (2.7) and (2.8), it then also follows that

$$\mu(n) > 0 \quad \text{for } n \in (\underline{n}, \bar{n}), \quad (2.9)$$

the proof of this being analogous to that which we later give for the corresponding result under poverty minimisation.

Equations (2.6) to (2.9) give rise to the few general qualitative conclusions available in the welfarist framework:

- (i) The marginal rate of tax should everywhere be non-negative;
- (ii) The marginal rate of tax on the lowest earner should be zero, so long as everyone supplies some labour at the optimum;³
- (iii) The marginal rate of tax on the highest earner should be zero, so long as wages in the population are bounded above.

Result (i) is perhaps more striking than is commonly recognised: while it may well be optimal for the average rate of tax on the least well-off to be negative, it cannot be desirable to subsidise their earnings at the margin. The limitations of the end-point results (ii) and (iii) are well known: simulations suggest that zero may be a bad approximation to optimal marginal tax rates in the tails of the distribution, and if it is optimal for some not

² For brevity, we ignore the possibility of a corner solution with hours worked equal to the time endowment.

³ More precisely, taking $n \rightarrow \underline{n}$ the requirement is that there be no bunching at zero hours (so that (2.6) holds at all wages above $\underline{n}$). Some additional restrictions on preferences and the wage distribution are also needed for the result: see Theorem 1 of Seade (1977).

to work then the optimal marginal tax rate at the bottom of the distribution can be shown to be strictly positive (Tuomala (1990)). Nevertheless, these results continue to colour professional thinking on issues of rate structure. The lower end-point result, in particular, has been taken as suggestive in arguing against the very high effective rates on the poor in the U.K. (as for instance by Kay and King (1986)).

Here then are three substantive implications of the welfarist approach. It is natural to ask whether they continue to apply when the objective is not social welfare maximisation but income poverty minimisation. Clearly if income here were to be interpreted as *equivalent* income in the sense of King (1983) then we have money metric utility and the welfarist conclusions will go through unaltered. But in policy discussions it is the minimisation of poverty defined on *measured* income - income actually received - which often appears to be the objective. The purpose of the next section is to see whether the welfarist conclusions survive in such a non-welfarist setting.

3. Income poverty minimisation: Theory

Suppose then that the sole aim of policy is to minimise an income-based poverty index of the general additively separable form analysed by Atkinson (1987):

$$P = \int_{\underline{n}}^{\bar{n}} G[x(n), x^*] f(n) dn \quad (3.1)$$

where $G[.]$ can be thought of as a generalised poverty gap: non-negative for $x \leq x^*$, zero otherwise, and satisfying

$$G_x[x, x^*] < 0 \text{ and } G_{xx}[x, x^*] > 0 \quad \forall x \in (0, x^*). \quad (3.2)$$

It is further assumed that

$$G_x[x, x^*] = 0. \quad (3.3)$$

This specification precludes a number of widely-used poverty indices, including the headcount ratio (which would have very strange implications in the present context). Nevertheless, the class of measures is a broad one.

As in the welfarist framework, individuals solve problem (2.3). The government faces the same revenue and incentive compatibility constraints. The only difference is that it now seeks to minimise poverty P . Denoting by λ and μ the multipliers on the two constraints, exactly as before, it is shown in the Appendix that the first order conditions imply, again where labour supply is strictly positive, that

$$t[z(n)] = \frac{G_x s}{\lambda} - \frac{\mu(n) u_x s_n}{\lambda f(n)} \quad (3.4)$$

where

$$\mu(n) = - \int_{\underline{n}}^n (G_x + \lambda) (1/u_x) \exp \left(- \int_p^n (u_{nx} / nu_x) dm \right) f(p) dp \quad (3.5)$$

satisfies

$$\mu(\underline{n}) = \mu(\bar{n}) = 0. \quad (3.6)$$

$$\mu(n) > 0 \quad \text{for } n \in (\underline{n}, \bar{n}). \quad (3.7)$$

To establish (3.7) note first, from (3.5), that the transversality conditions (3.6) imply that $G_x + \lambda$ changes sign as n increases. Since G is strictly convex (by assumption) and $x(n)$ is increasing in n (as a consequence of the preference restriction (2.2)), $G_x + \lambda$ is strictly increasing in n . Thus $-(G_x + \lambda)$ must change sign only once, and must start off positive. But then if μ were ever to become negative it could not recover to $\mu(\bar{n})=0$.

The characterisation (3.4) has an appealing interpretation. Recalling (2.6), and comparing (3.5) with (2.7), note that the second term in (3.4) is of precisely the same form as the welfarist tax formula except that $W'u_x$, the social marginal valuation of consumption, is replaced by $-G_x$, the marginal reduction in the generalised poverty gap associated with an increase in consumption. Loosely speaking, this term can thus be thought of as corresponding to a welfarist calculation based on a social valuation of the form $G[v(n), x^*]$; measuring poverty, that is, in terms of shortfalls of utility from some threshold. But our concern is with income poverty, so that in this sense it is the first term in (3.4) that captures the distinctive features of the non-welfarist approach. And the direction in which the concern with income rather than utility points then emerges as unambiguous: it tends towards lower marginal tax rates on the working poor. The reason is straightforward. Consider the local effect of a compensated reduction in the

marginal tax rate at some point in the schedule. This will lead to an increase in the gross incomes of those located at that point, associated with which will be an increase in their net incomes of $-u_y/nu_x = s$. From the welfarist perspective, these consequences for individual behaviour are a matter of indifference, since only utility matters and that is unchanged. In terms of the present non-welfaristic perspective, however, the higher net income induced by this lowering of the marginal tax rate is desirable in itself: it reduces the generalised poverty gap by the amount $G_x s$ appearing in the first term of (3.4).

Do the three central welfarist results continue to apply? Result (iii) clearly does. Using $\mu(\bar{n}) = 0$ in (3.4) and noting that $G_x[x(\bar{n}), x^*] = 0$ (so long as the highest earner is not poor), it follows that $t[z(\bar{n})] = 0$: the marginal rate at the top of a bounded distribution should again be zero. This is as one would expect. For in the context of poverty alleviation the only reason to care about the highest earner - indeed about any of the non-poor - is as a source of revenue, and it is well known that in these circumstances one would want a zero marginal rate at the top: if it were strictly positive, additional revenue could be extracted by slightly lowering it and thereby inducing the highest earner to earn additional taxed income.

Results (i) and (ii), in contrast, are overturned. Taking limits in (3.4) - thereby assuming that it is optimal for all to work - and using the first part of (3.6) one finds

$$t(\underline{n}) = G_x[x(\underline{n}), x^*]s[x(\underline{n}), z(\underline{n}), \underline{n}]/\lambda < 0 \quad (3.8)$$

by the first condition in (3.2) (and assuming too that poverty is not entirely eliminated). In these circumstances the marginal rate at the bottom of the gross income distribution should be strictly negative: if it is optimal to have everybody work, poverty alleviation calls for a marginal subsidy on the earnings of the very poorest. Indeed the conclusion to be drawn here is of a rather more general kind than the welfarist lower end-point result (ii) above: whereas the latter applies only at the lower extreme of the distribution, a set of measure zero, (3.8) implies - given continuity - that there exists an interval over which a negative marginal rate is appropriate.

The rationale for this result can be seen from Figure 1. Suppose that the initial position is one in which the marginal tax rate on the very poorest household is indeed strictly negative. This initial equilibrium is shown as point α , an indifference curve of an n -household being tangential to the segment AA of the budget constraint implied by the tax system in force. (For clarity, but inessentially, budget constraints are drawn as if they were linear in the relevant range). The assumption of a negative marginal rate

implies that the slope of AA exceeds that of CC, which is the 45° line through α . Consider now a tax reform that increases the marginal rate at α whilst retaining α itself as a feasible point; diagrammatically, the budget constraint rotates clockwise about α to arrive at BB. Though not shown, one can imagine that the tax function is simultaneously amended further up the distribution so as to leave all other households unaffected. Is this reform desirable? From the perspective of social welfare maximisation it certainly is, there then being two effects pointing in the same direction. The first is that this poorest household now attains a higher utility level, moving to a point like β . The second is that tax revenue increases: since β lies below CC

net income falls by more than gross. From the perspective of poverty alleviation, however, judgement must be suspended. For while the increase in revenue is still to be welcomed it is then income that matters, not utility; and the labour supply response to this reform has the effect of reducing the poorest household's net income, so deepening its measured poverty. It is the need to balance considerations of this kind that gives rise - when it is optimal to induce all to work - to negative marginal rates on the lowest earners.

What though if it is optimal for the poorest individuals to be idle? There are then two cases to consider. The first is that in which $\underline{n} > 0$, so that everyone is capable of earning at least some small income. It can then be optimal for some not to work only if the marginal rate is non-positive at the bottom of the distribution. For suppose, to the contrary, that the marginal rate at the bottom is strictly positive and that the poorest individuals work zero hours. Let z' be some level of gross income that lies within the range over which the marginal rate is strictly positive and which is also less than $\underline{n}L$, where L denotes the time endowment. Now consider a reform which leaves the tax schedule at and above z' unchanged while imposing an average rate of 100% below z' . Those who had previously earned less than z' , including in particular the idle, will now choose to earn exactly z' ; all others will be unaffected. The consumption associated with z' can be no less than that previously associated with lower gross incomes, otherwise those initially at z' would have chosen to work fewer hours. Therefore poverty cannot be increased by this reform. Strict positivity of the marginal rate up to z' implies, moreover, that the labour supply effects of the reform lead to higher tax revenue. And with this increased revenue it will generally be possible to bring about a strict reduction in poverty. Thus the initial schedule cannot be optimal. The second and less clear-cut case is that in which $\underline{n} = 0$. The preceding argument then fails, since for any strictly positive level of gross income there will be a measurable set of individuals unable to earn that income even by working all of their time endowment. Tax reforms of the kind just described, intended to induce those who can earn some income to do so and thereby reduce the revenue cost of their support, will inescapably reduce the consumption of the very poorest and to that extent worsen poverty. The desirability of such a reform then becomes

unclear, being liable to depend, for instance, on the form of the poverty index. We have been unable to find any general results on the sign of the optimal marginal rate at the bottom of the distribution when $\underline{n} = 0$.

The possibility of optimally negative marginal tax rates is certainly confined, however, to the poorest of the poor. At the poverty line wage n^* , defined by $x(n^*) = x^*$, one finds (from (3.3), (3.4) and (3.7)) that

$$t(n^*) > 0. \quad (3.9)$$

Thus, invoking continuity once more, there exists some wage strictly below that required to escape poverty such that all those with higher wages - whether poor or not - face a positive marginal tax rate.

The results of this section thus point towards a pattern of marginal tax rates below the poverty line that is both complex and potentially of a kind very different from that to which the welfarist tradition has pointed. But there remain the questions of how far the considerations that point to low or even negative marginal tax rates on the very poorest individuals extend into the range of incomes, and of how the poverty-minimising rate structure is affected by the precise location of the poverty line x^* and by the form of the poverty gap measure $G[.]$. These issues can be addressed only by numerical simulation, a task to which we now turn.

4. Numerical simulations

Recognising both the difficulties of obtaining closed form solutions for optimal non-linear taxes and the potential limitations of end-point results, the welfarist literature has developed a tradition of numerical simulation (a tradition reviewed in Tuomala (1990)). The original Mirrlees simulations, and many others since, have been carried out for the benchmark case of Cobb-Douglas preferences

$$u = \ln(x) + \ln(1-y) \quad (4.1)$$

(the time endowment being normalised at unity) and under the assumption that n is distributed lognormally, with the mean of $\ln(n)$ being -1 and its

standard deviation 0.39. We retain these basic specifications. The revenue requirement R will be taken to be 10% of gross income, again a conventional figure (intended as a very rough approximation to the levels of expenditure on public goods commonly observed). The essential novelty is the form of the objective function. For this we take a poverty index of the form developed by Foster, Greer and Thorbecke (1984):

$$P^\alpha = \int_{\underline{n}}^{\underline{n}^*} \left(\frac{x(n) - x^*}{x^*} \right)^\alpha f(n) dn, \quad \alpha > 1. \quad (4.2)$$

This has been widely used in the analytical literature on targeting (as for instance in Besley (1990) and Kanbur and Keen (1989)), the parameter α providing a convenient parameterisation of alternative degrees of poverty aversion.

One implication of this specification should be noted. With Cobb-Douglas preferences (so that the marginal rate of substitution between consumption and work is strictly positive at zero hours) and a lognormal wage distribution (so that $\underline{n} = 0$), there are some who will work only if the marginal tax rate at the bottom of the distribution is infinitely negative. In both the welfarist context and that of poverty minimisation one would then expect it to be optimal to have some of the population idle. In the welfarist case, the optimal marginal rate at the bottom must then be strictly positive. As discussed in the preceding section, however, when the objective is to minimise income poverty and $\underline{n} = 0$ the sign of the optimal marginal rate at the lower end-point is in general ambiguous. The simulations can thus provide some indication of the extent to which the argument for non-positive marginal rates at the lower end when $\underline{n} > 0$ continues to exert some force when instead the wage distribution is not bounded away from zero.

Simulation results are reported in Table 1 below, which gives optimal average and marginal tax rates at various percentiles of the wage distribution, starting with the marginal tax rate at the bottom and including the point at which the assumed poverty line is to be found. Panels (a)-(c) all take $\alpha=2$, and differ in taking successively higher poverty lines. Panel (d) looks at

Table 1. Simulation results

Panel (a): $F(n^*)=0.31$, $\alpha=2$, $X/Z=0.9$

	$F(n)$	ATR(%)	MTR(%)
x^*	0.06	-100	69
	0.31	-3	62
	0.50	12	53
	0.90	29	35
	0.99	29	23
$F(n_0) = 0.06$, $x(n_0) = 0.06$.			

Panel (b): $F(n^*)=0.43$, $\alpha=2$, $X/Z=0.9$

	$F(n)$	ATR(%)	MTR(%)
x^*	0.02	-100	63
	0.43	0	54
	0.50	9	53
	0.90	27	34
	0.99	27	17
$F(n_0) = 0.02$, $x(n_0) = 0.06$.			

Panel (c): $F(n^*)=0.56$, $\alpha=2$, $X/Z=0.9$

	$F(n)$	ATR(%)	MTR(%)
x^*	0.003	-87	56
	0.50	8	54
	0.56	16	48
	0.90	26	34
	0.99	26	17
$F(n_0) = 0.003$, $x(n_0) = 0.06$			

Panel (d): Maximin ($\alpha=\infty$), $X/Z=0.9$

$F(n)$	ATR(%)	MTR(%)
0.16	-100	73
0.50	17	53
0.90	32	35
0.99	32	26
$F(n_0) = 0.16, x(n_0) = 0.07$		

Notes: (a) n_0 denotes the wage below which individuals choose not to work.

(b) X/Z denotes the ratio of aggregate consumption to aggregate output.

(c) ATR (MTR) denotes the average (marginal) tax rate.

the maximin case,⁴ which corresponds to $\alpha=\infty$.

Several features of these results are noteworthy. First, and as anticipated, it is always optimal to have some of the poorest individuals idle. Second, the marginal rate on the lowest gross income - which, as just noted, is ambiguous in sign when $\underline{n} = 0$ - emerges as very strongly positive: not only is it not negative, it is not even low. Third, marginal tax rates decline monotonically from the poorest to the richest individuals. This is in marked contrast to the welfarist tradition, in which the pattern of optimal marginal rates typically has a mild inverse-U shape. Indeed it is notable that the poverty-minimising schedules - unlike those of Mirrlees (1971)- could not be described as approximately linear. This confirms our earlier remarks on the importance of relaxing non-linearity, and, moreover, tends to weaken the common argument that - given its administrative advantages - the social loss through restricting oneself to linear income taxes is likely to be small. Fourth, comparing panels (a) to (c), increases in the poverty line lead to a

⁴Since the poorest individuals do not work at the optimum for either the welfarist or poverty minimisation problems, 'maximin' here refers simultaneously to both income and utility.

reduction in optimal marginal rates at and below the poverty line. Intuitively, the explanation for this appears to be that the case for low marginal rates in order to encourage those at or near the poverty line to move over it becomes stronger as the poverty line moves into denser parts of the distribution. Fifth, comparing panel (d) with the rest, increases in the extent of aversion to inequality tend to increase marginal rates on the poor. This is perhaps as would be expected, since the greater one's concern with poverty alleviation the more attractive schemes approaching minimum income guarantees are likely to be.

But perhaps the most important feature of the results is the finding of marginal tax rates on the poor that are invariably rather high (bearing in mind the fairly minimal revenue requirement). In most cases marginal rates on the bulk of the poor exceed 60%, and in all cases they exceed 50%. These rates are somewhat lower than those (of around 80 to 90%) found by Garfinkel, Moreland and Sadka (1982) for the welfarist case. To that extent the informal argument following equation (3.4) that the non-welfarist perspective is likely to point towards somewhat lower marginal rates on the working poor is borne out. Nevertheless, the case for low marginal tax rates fails to leave as discernible a trace in the simulations as one might have expected. Even with the relatively elastic labour supply responses implicit in Cobb-Douglas preferences, a stronger mark is left by the case for high marginal rates associated with the unattainable ideal of perfect targeting described at the outset.

5. Conclusion

The central thrust of the present non-welfarist analysis is to strengthen the case for lower marginal tax rates on the poor. But not by much. This alternative perspective has been shown to overturn the lower end-point result of the traditional literature: when everybody works at the optimum, the minimisation of income poverty requires that the lowest earner face a strictly *negative* marginal tax rate and, by continuity, that the marginal rates on the not-quite-so-poor also be negative. However, we have also seen that the optimal marginal tax rate at the poverty line is strictly positive.

Putting these two observations together points to a potentially complex rate structure, the shape of which can only be fleshed out by numerical methods. These suggest optimal marginal rates on the poor that are only a little lower than is typical in a welfarist framework. But while the rates we calculate - unlike those we observe - are nowhere near 100%, they are not close to zero either. Taking conventional functional forms and parameter values, the income poverty minimising marginal tax rates on the poor are in the order of 60-70%.

APPENDIX

A. Derivation of (3.4) and (3.5)

Substituting for x from the budget constraint, differentiation of $u(x(n), y(n))$ with respect to n gives

$$du/dn = u_x(1-t)y + \{u_x(1-t)n + u_y\}dy/dn. \quad (A.1)$$

Using the first order condition (2.3), individual optimisation thus implies the envelope condition

$$du/dn = -yu_y/n \equiv u_n(x, y, n). \quad (A.2)$$

To simplify the optimisation, we take u and y to be the objects of choice. Inverting direct utility then gives $x = h(u, y)$, where

$$h_y = -u_y/u_x, \quad h_u = 1/u_x. \quad (A.3)$$

Defining too $g(u, y, n) = u_n[h(u, y), y, n]$, it is straightforward to check that

$$g_y = -nu_x s_n, \quad g_u = u_{nx}/u_x. \quad (A.4)$$

The optimum is then characterised by a pair of functions $u(n)$ and $y(n)$ which minimise the poverty index subject to the incentive compatibility condition $du/dn = g$ and the overall resource constraint. Introducing multipliers λ and $\mu(n)$ for these constraints and integrating by parts, the Lagrangean becomes

$$L = \int_{\underline{n}}^{\bar{n}} \left([-G(x, x^*) + \lambda(ny - x)]f(n) - \mu' u - \mu g \right) dn + \mu(\bar{n})u(\bar{n}) - \mu(\underline{n})u(\underline{n}). \quad (A.5)$$

Differentiating with respect to u and y gives the first order conditions:

$$L_u = -[G_x + \lambda]h_u f(n) - \mu' - (\mu(n)u_{nx}/u_x) = 0 \quad (A.6)$$

$$L_y = [-G_{xy}h_y + \lambda(n-h_y)]f(n) + \mu(n)nu_x s_n = 0. \quad (A.7)$$

Dividing (A.7) by λf and rearranging, one finds

$$1 - (h_y/n) = (G_{xy}h_y/\lambda n) - (\mu u_x s_n/\lambda f). \quad (A.8)$$

Using (2.3) and the first part of (A.3), (A.8) becomes (3.4). Solving (A.6), using (3.6), gives (3.5).

B. Outline of the computation method

For an additively separable utility function, as in (4.1), $u_{nx} = 0$ and so (A.6) becomes

$$-[G_x + \lambda]f(n)/u_x - \mu' = 0. \quad (B.1)$$

Using (A.9), the solution of (B.1) is

$$\mu(n) = -\int_{\underline{n}}^n (G_x + \lambda)(1/u_x)f(p)dp. \quad (B.2)$$

Equation (A.7) is now

$$[-G_{xy}h_y + \lambda(n-h_y)]f(n) + \mu(n)(u_y + yu_{yy})/n = 0 \quad (B.3)$$

which can be written as

$$n^2 r f(n) - \int_{\underline{n}}^n [G_x + \lambda](1/\lambda u_x)f(p)dp = 0 \quad (B.4)$$

where $r = [(G_{xy}u_y/\lambda n u_x) + 1 + (u_y/n u_x)]/(u_y + yu_{yy})$. Differentiating (B.4) gives the differential equation

$$dr/dn = \begin{cases} -r[2 + (nf'/f)]/n + [G_x/\lambda u_x] + (1/u_x)/n^2 & ; x < x^* \\ -r[2 + (nf'/f)]/n + 1/u_x n^2 & ; x > x^*. \end{cases} \quad (B.5)$$

The formulae (B.5) and (A.2) form a pair of non-linear differential equations in u and r . They provide the solution to our problem, together with the conditions (2.5) and $\mu(\bar{n}) = 0$ ($rn^2 f \rightarrow 0$ as $n \rightarrow \infty$).

A fourth order Runge-Kutta method is used to solve the non-linear differential equation system (B.5) and (A.2). Instead of solving for λ , a value is assumed for it. It can be shown that there is a critical n_0 such that

$$\begin{aligned} y(n) &= 0 \text{ for } n \leq n_0 \\ y(n) &> 0 \text{ for } n > n_0. \end{aligned} \quad (B.6)$$

Trial and error is used to find a value of n_0 which enables us to satisfy the condition $y(n_0) = 0$. From (B.4) we can solve $x(n_0)$ using the Newton method. When λ is given and $x(n_0)$ is solved, the integration of (B.5) and (A.2) can be started. When u^{i+1} and r^{i+1} are obtained (i referring to an iteration cycle), the new values for y and x are calculated from (B.5) and the utility function by using the Newton method.

References

- Atkinson, A.B. (1987). 'On the measurement of poverty', Econometrica 55, 749-764.
- Besley, T. (1990). 'Means testing versus universal provision in poverty alleviation programmes', Economica 57, 119-129.
- Besley, T. and R. Kanbur (1988). 'Food subsidies and poverty alleviation', Economic Journal 92, 701-719.
- Dilnot, A. and G. Stark (1989). 'The poverty trap, tax cuts and the reform of social security'. In Dilnot and Walker (1989).
- Dilnot, A., and I. Walker (eds) (1989). The Economics of Social Security (Clarendon Press, Oxford).
- Foster, J., J. Greer and E. Thorbecke (1984). 'A class of decomposable poverty measures', Econometrica 52, 761-766.
- Garfinkel, I., K. Moreland and E. Sadka (1982). 'Income testing and social welfare: An optimal tax-transfer model'. In I. Garfinkel (ed) Income Tested Transfer Programs: The Case For and Against (Academic Press, New York).
- Kanbur, R. and M.J. Keen (1989). 'Poverty, incentives and linear income taxation'. In Dilnot and Walker (1989).
- Kay, J.A. and M.A. King (1986). The British Tax System, 4th edition (Oxford University Press, Oxford).
- King, M.A. (1983). 'Welfare analysis of tax reforms using household data', Journal of Public Economics 21, 183-215.
- Mirrlees, J.A. (1971). 'An exploration in the theory of optimum income taxation'. Review of Economic Studies 38, 175-208.
- Seade, J. (1977). 'On the shape of optimal tax schedules'. Journal of Public Economics 7, 203-235.
- Seade, J. (1980). 'Optimal non-linear policies for non-utilitarian motives'. In Collard, Lecomber and Slater (eds) Income Distribution: The Limits to Redistribution (John Wright and Sons, Bristol).
- Seade, J. (1982). 'On the sign of the optimum marginal income tax'. Review of Economic Studies 49, 637-643.
- Tuomala, M. (1990). Optimal Income Tax and Redistribution (Clarendon Press, Oxford).