

ON THE SHAPE OF OPTIMAL TAX SCHEDULES

J. K. SEADE*

University of Warwick, Coventry CV4 7AL, England

Received January 1976, revised version received June 1976

Take consumers to be described by a parameter h (skill, needs, etc.) with utilities defined on N commodities, including factor supplies. Our main result is that if in the optimum each component of this consumption vector is bounded and bounded away from zero over the population, each marginal tax must be zero at both ends of the corresponding tax schedule. For the income tax case, it is also shown that if the marginal tax rate at the top of the scale is positive, one can construct another tax schedule which is strongly Pareto superior (dominating the first one at all or most income levels), requirements of information being low.

1. Introduction

An old question in the normative theory of taxation is how best to decide on the trade-off between efficiency and equality. Pigou (1946) stressed that since the (social) marginal utility of income decreases with income, any transfer of purchasing power from rich to poor which does not decrease aggregate output is socially desirable. However, if we can only achieve this redistribution through taxation (and assuming away undistorting, once-and-for-all reallocations of property), it becomes central to consider to what extent the supply of efforts, and hence output, would be affected if a given redistributive tax-scheme were imposed.¹ Specifically, we want to find a way to take this efficiency-effect into account when selecting the progressivity of tax schedules. This is the central feature of the theory of optimal income taxation, originated by Mirrlees (1971), to incorporate the analysis of the different supplies of labour brought about by alternative tax schedules into a welfare maximization formulation.

The immediate aim when setting up an optimization exercise of this sort, apart from the basic one of providing a rigorous and workable formulation of the problem, clearly is to characterize an optimum fully by means of necessary conditions that it must satisfy. These conditions, however, are not of so much

*This paper is part of a D.Phil. thesis (Oxford), for which Professor J. A. Mirrlees provided most helpful supervision. My special thanks are due to him. I am also indebted to Professor A.B. Atkinson, J.S. Flemming, who also supervised part of my work, and Dr. N.H. Stern, for very valuable suggestions, as well as to the members of seminars at Maryland, Warwick and Nuffield College, Oxford for their useful comments. The paper was presented at the 1976 European Meeting of the Econometric Society (Helsinki). Financial support from the Ford Foundation is gratefully acknowledged, as is assistance from CONACYT (Mexico).

¹Sidgwick was already aware of the likely negative incentive effect of an increased equalization of incomes [see reference in Atkinson (1973, p. 92)].

interest in themselves, but rather as a source of more specific knowledge of the qualitative nature and quantitative features of optimal taxes. By and large, the dominant line of approach to this problem has been the numerical analysis of solutions for various specific examples. Much is already known, for instance, on the way the general *level* of taxes depends on the specification of the model, notably that optimal marginal taxes are surprisingly low in many cases, but that they tend to increase: (i) the more pre-tax inequality there is [Mirrlees (1971)], (ii) the more inequality-averse the government is [Atkinson (1973)] and, notably, (iii) the less a typical consumer's supply of labour reacts to small (and compensated) changes in his marginal tax rate – a result reminiscent of the Pigovian claim mentioned above [Feldstein (1973), Stern (1976)]. Another, perhaps less intuitive feature of solutions, relates to the *shape* of tax schedules: does the analysis of optimal taxation provide support for the general presumption that the taxation of income should be progressive² throughout? Surprisingly, numerical examples in most cases give a negative answer to this question, with optimal income tax schedules typically being progressive at low levels of income and then regressive from a certain point onwards.³ The incidence of this feature invites an intuitive or a formal explanation.

In contrast to this relatively large amount of mainly numerical research, few analytical results have been found (except for particular cases), which is to say that in fact not much is known of the qualitative nature of optimal income taxes in the general case. I would say there are only two results in this direction: that optimal marginal income taxes always lie between zero and one [Mirrlees (1971, propositions 2 and 3)]⁴ and that the marginal tax at the top of a bounded optimal income tax schedule must be zero⁵ – providing a firm rejection of uniform progressivity of optimal taxes (for bounded incomes). It is in more general results

²Two distinct concepts of progressivity of a tax are commonly met, corresponding to increases of average or of marginal tax payments with income. The former corresponds more closely with the intuitive idea of 'progressivity' as a measure of how redistributive the tax is, but the two concepts clearly are closely interwound. I here take progressivity in the marginal tax sense, although the implications of our remarks and results for the shape of average taxes can easily be worked out. 'Regressivity' is similarly defined, as the converse of progressivity.

³Two main exceptions to this general shape have been noticed: one is that with a Pareto distribution of skills [for which Mirrlees (1971) found rather high asymptotic marginal taxes] and in the maximin case (presumably in more general cases too) the optimal marginal tax rate is *monotonically* increasing [Atkinson (1972)]. This is related to the discussion in footnote 30 (below) and the text preceding it. Another exception, easier to rationalize, is when the supply of labour is positively sloped at all high levels of (net) income [Mirrlees (1971, case (i) on p. 15)], in which case not much money needs to be left to the rich from their marginal earnings to have them always working more the more able they are, thus making optimal marginal tax rates tend asymptotically to one, so as to exploit these rich revenue (or redistributive) possibilities.

⁴Mirrlees also found that in the optimum arrangement there normally is a fraction of the population (all those below a certain ability level) who remain idle, but this result needs to be qualified (section 3, below).

⁵See Phelps (1973) for the maximin criterion and Sadka (1976) for the utilitarian case. The 'top' of a tax curve is taken to be at the highest level of income actually observed in the population, although the tax schedule may of course be defined for *all* levels of income.

of this type that I am interested here. Important though this result is, if only because it casts some doubt on the desirability of tax schedules observed in practice, it however provides no hint as to how the rest of the tax schedule ought to be. In particular, it is not unlikely that it be precisely the tax treatment of the *bottom* fractions of the population, not of those at the top, that has more bearing on welfare, both because there usually are more people near the bottom than near the top of the scale of incomes and because each one of them has a higher welfare weight. It is therefore of particular interest to try to arrive at tax rules of the same degree of generality for the bottom end of the income tax. Since a zero marginal tax at the top implies regressivity near the top, should we then have uniform regressivity in some cases? Do we want high or low marginal taxes near the bottom?

The main purpose of this paper is to show that the zero marginal tax is a quite general property of *both* endpoints of optimal nonlinear taxes (i.e. those whose marginal rate of taxation is not restricted to constancy), such as taxes on income, expenditure or wealth. Specifically, in the particular case of the income tax, we find that the zero-marginal-tax result we know for the top of the scale also holds at the bottom, as long as the economy has no zero-income households in the optimal arrangement. More generally, marginal taxes at both ends of any optimal nonlinear tax schedule are zero, the general requirement being that the consumption vector to which the tax gives rise (including the amount of work done and hence gross income) be bounded above and away from zero across the population in the optimum, a not unlikely condition when we are dealing with aggregates such as income, total consumption or expenditure in each period of a man's life. This result is presented in section 5 in this fairly general form, applying to unspecified nonlinear taxes or tax mixes. The model required is formulated in section 4 and the discussion of the result is concentrated, mainly, in section 6.

But before we go on to the main sections, it will be convenient to consider in particular some aspects of the income tax case. In section 2 (and the appendix) we revise the familiar zero-marginal-tax result for the top-end of income taxes with the purpose of extending it in directions not covered by the otherwise more general formulation of later sections, also presenting a way to *improve* on existing suboptimal income taxes that has some practical appeal, and in section 3 we discuss and qualify the well-known result that there should be a number of people not doing any work in the optimal tax regime, which has direct bearing on whether the marginal tax at the bottom of the income scales should be zero or not.

2. The zero marginal income tax at the top

2.1. A necessary condition for optimality

The examples studied by Mirrlees and Atkinson had populations with

unbounded individual income (or wage rates, more precisely), that is, with upper tails of incomes extending indefinitely to the right, which is a feature of most density functions commonly used. Although in these examples the marginal tax rate did not always tend to zero, it did often show a tendency to decline -- in some cases in fact towards zero. One may conjecture that the force behind this asymptotic behaviour of the tax rate is revealed by the bounded income tax case, i.e. that in which we assume that no one in the population has a wage rate higher than a certain level, for in this case a clear-cut result is available, that the marginal tax should be zero at the top of the scale [Phelps (1973), Sadka (1976)]. The intuitive explanation of this result is immediate in the utilitarian case: suppose we have any given tax schedule whose marginal tax at the top is not zero but positive.⁶ Allow then the top man to earn tax free any excess of income he had previously decided not to make the effort to earn, partly because he would have been paying tax on it. In taking advantage of the change he ends up better off, working (and consuming) more and paying no less tax than before. The same is true of any other people who may decide to work harder so as to take advantage of the new tax-free range at the top. Hence the change is good and is revenue-feasible.⁷

Geometrically (see fig. 1), denote by $c \equiv c(w)$ and $l \equiv l(w)$ the consumption and hours worked by a man whose wage rate is w and whose income is $y \equiv y(w) = wl(w)$. We assume perfect substitutability in production between labour of different grades, 1 hour of a w -man being identical to w hours of a 1-man. Let θ , drawn in the diagram as $A\theta$, be the net-income function imposed by the government and from which each person must make a choice. (θ will be used both as a label to denote a particular net income curve and as the corresponding functional relation.) We assume that all consumers have the same indifference curves (not shown) between consumption c and hours worked l , from which we can draw their indifference curves between c and wl , i.e. between net income (\equiv consumption) and gross income. These curves are obtained simply by expanding or contracting horizontally each (c, l) -curve by a factor w .⁸ Now suppose that at the point a selected by the top man, the marginal tax rate is positive, i.e. $\theta'(y) < 1$ at that point (because net income is related to tax payments $\tau(y)$ by $\theta(y) = y - \tau(y)$, hence $\theta' = 1 - \tau'$). It is then clear that θ cannot be optimal, at least as long as one attaches positive weights to the utilities of the highest earners, because if we replace the segment $a\theta$ by any steeper one such as $a\hat{\theta}$ (which is a zero-marginal-tax line), the top man moves to some point

⁶The above-mentioned result by Mirrlees (1971) that no marginal taxes can be negative in the optimum is central to the argument.

⁷I partially followed Stern's (1976) way of putting the result verbally.

⁸Hence each w has a different map of indifference curves (c, y) , reflecting the different amounts of work different people would have to do to earn a given pre-tax income. If (but not only if) c is noninferior these curves are at each point strictly flatter (lower $|MRS_{cy}|$) the larger w is, implying that income is non-decreasing in the wage rate and strictly increasing if (or wherever) the tax schedule is smooth.

a^* where he is better off; it is not hard to see that there will normally be other men (those close enough in ability to the top man, if there are any) deciding now to work harder so as to earn income in the stretch aa^* who will be better off as well and, finally, the change is feasible, for all these movements are along lines with slope less than (or equal to) one, i.e. with each consumer whose behaviour is affected by the tax change increasing his gross income by more than (or as much as) his consumption, therefore not decreasing his tax payments. Meanwhile, the rest of the population remain earning, consuming and paying taxes in the same amounts as before the tax change. Hence the well-known result that θ cannot be optimal.

It is worth noting at this point an immediate yet interesting generalization of this result: although the theory of optimal taxation becomes considerably more complicated once we allow for people to differ in more than one respect [see Mirrlees (1976), Seade (1977)], e.g. wage rate and work preferences, it still is a necessary condition for an optimum that the marginal tax faced by the highest earner should be zero. In the simple model with identical preferences, all consumers earning the same level of income also have the same wage rate (apart from possible corners of θ) and hence all the indifference curves tangent to θ at each point are equal. When preferences vary in unspecified ways across consumers, this will not be the case any more; some people will be earning a given income because they are more able, while others because they are more hard-working, say. Similarly, the top-income man will now not necessarily be the one with the highest wage rate (or natural skill, rather, for we can allow for strength of efforts to be a choice variable, which of course affects the wage rate). But all this does not really affect our conclusions here, because it is still true that if we replace the segment $a\theta$ in fig. 1 by $a\hat{\theta}$, the new net income schedule will cut the initially-optimal indifference curves of some people at or near the top, allowing them to move towards higher levels of preference.⁹ Anyone who for this reason decides to work more (or harder) will be better off and will be paying no less tax than before. Thus, the optimal income-tax schedule must have a zero marginal tax at the top, even when the preference structures underlying the work-behaviour of different consumers differ in any number of ways.

Notice that the result imposes regressivity of optimal income taxes near the top, but perhaps very locally: nothing is implied on whether the tax schedule as a whole should be very redistributive or not. Nevertheless, it suffices to establish not only suboptimality but indeed Pareto-inefficiency of income taxes with a positive marginal tax at the top. As an analytical proposition it is fairly strong. However, one would like to go beyond that and seek some applicable policy

⁹Certainly all those near enough the top and with smooth indifference curves (although smoothness is not really necessary) and certainly no one with Leontief-type indifference curves, who would prefer to stay where they were. The final ordering of incomes will now normally change, unlike the simple case.

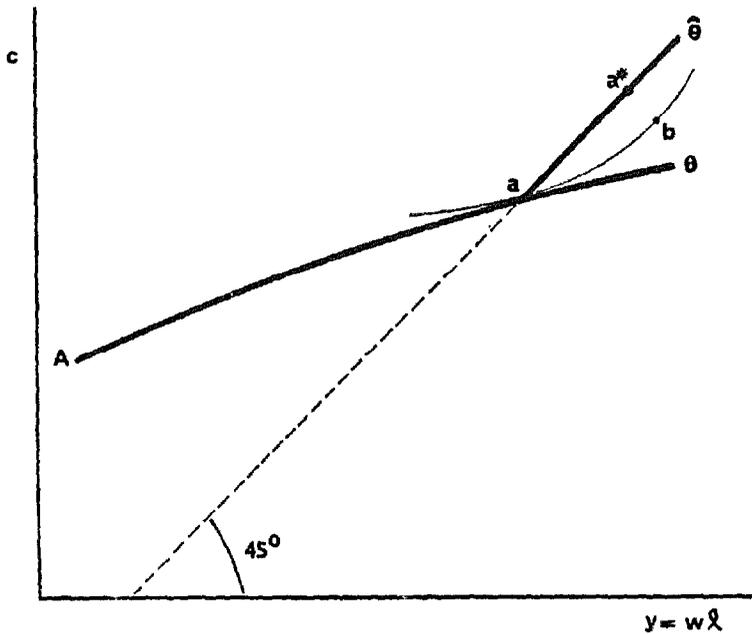


Fig. 1. Suboptimality of θ .

implications from the analysis, and in this direction this simple tax reform does not seem to have much appeal, as it improves matters by raising the well-being of a very small group of people who are, moreover, those at the very top. True, the change will probably yield some extra revenue which can be reallocated somewhere else in the economy for the benefit of others. However, the first-order utility benefit to the rich will yield a second-order extra amount of revenue and hence a second order welfare benefit to the rest of the population. This is clearly illustrated taking the extreme case of a discrete population distributed near the top in such a way that it does not pay the second man to move onto the tax-free segment aa^* , in which case the only beneficiary of the change is the top man, who pays no extra revenue at all.

We can alternatively devise a tax reform at the top specifically aimed at increasing the tax payments by the direct beneficiaries of the tax change, including the top man: instead of reducing the marginal tax immediately to zero at the level of income the top man initially had, we can reduce it first by a lesser amount and only at some higher point perhaps make it zero, so that the new tax schedule is one that lies somewhere between θ and $\hat{\theta}$ in fig. 1. Under such a scheme it is easy to show that the tax payments by anyone who decides to take advantage of the tax change by working more are always higher than under the tax schedule $Aa\hat{\theta}$ but that, in exchange, there are fewer such people (this can easily be checked drawing in fig. 1 the relevant indifference curve of the man who is just indifferent between moving onto $\hat{\theta}$ or staying where he was on θ , who definitely prefers to stay on θ if faced with a tax change at the top that lies strictly

below $\hat{\theta}$). The extreme case is when the new tax schedule leaves the top man at point **b** (fig. 1) on his initial indifference curve, so that as much revenue as possible is being exacted from him (within the class of Pareto-improved schedules). In this case the relative orders of magnitude of the top man's utility gains and tax-payments changes have been completely reversed from what they were in the simple tax change of fig. 1, but now no one in the economy other than the top man reacts to the tax change at all.

2.2. *Nonlocal Pareto-improvements*

The tax reform outlined in fig. 1 and its variants just mentioned have the common feature that they leave the current tax schedule up to its current top-point basically unchanged, introducing certain desirable changes only from that point onwards (and, indirectly, shifting the whole tax schedule slightly upwards, if the increased revenue is redistributed as a poll sum). The result of this is that whether we want to increase the utilities of the very rich or their tax payments, we are in any case relying only on the reactions of that group to improved incentives. In this sense the result is a very local one, which reduces its practical appeal.¹⁰ It would be of more interest to find an improvement on a given inefficient tax schedule such that, on the one hand, the benefit is shared by the population at large and, on the other, the source of this benefit is not a local change of behaviour but one of the population at large too, a feasible upwards shift of all of θ . Let us state the following result that says that this is possible, proceeding then to prove it by constructing the required variation of θ . I return to the standard assumption of equal indifference maps between consumption and hours worked for all consumers. I also assume consumption is not inferior, so that income is nondecreasing in the wage rate (see footnote 8) and, only for simplicity, that the net-income function θ has no corners.

Theorem 1. Let income be bounded in the population. Then any income-tax schedule τ with a positive marginal tax at the top of the scale can be replaced by another one that leaves all consumers better off (or all those with positive incomes and above the highest level at which the marginal tax is zero, if the original schedule τ has $\tau'(y) = 0$ anywhere), inducing them to earn more income while paying the same amount of tax as before.

¹⁰It also presents some practical difficulties: if the increased revenue collected at the top is to be distributed elsewhere in the scale, say as a poll grant, the diagrammatic neatness of the analysis is lost, first because so as to know how much revenue is collected at the top and hence how big the poll transfer can be, one needs to know the distribution of skills in the population; and second, because recipients of the increased poll grant will themselves adjust (normally reduce) their supply of labour and hence their tax payments, thus requiring a more elaborate analysis.

Hence the original tax schedule is strongly Pareto-inferior and of course suboptimal.

To prove this, let θ in fig. 2 be the initial net consumption curve, to which a tax schedule $\tau = y - \theta$ is uniquely associated. We draw indifference curves for each w (constructed from the unique map on (c, l) and the different values of w) which, when tangent to θ , show the gross-income-net-income point a w -man selects, just as before. Notice that any movement of a person's income and consumption along the 45° line on which he is initially located represents equal total tax payments, that is, such movements are feasible from the point of view of budget balance. Equivalently, assuming labour is paid its marginal product, these movements are also feasible from the point of view of total output, since along these lines each consumer's consumption and output (i.e. income) change by the same amounts.¹¹ Consider now the indifference curves of a w -man, who facing the net income schedule θ chooses to earn income and consume at point P in the figure. Let us assume that the marginal tax he faces is positive (see footnote 6), i.e. that the slope of θ and hence the slope of his indifference curve through P are less than one. It is clear that if this man were allowed to choose his most preferred position on the 45° line that passes through P he would choose a point of more income (and work), i.e. one that lies northeast of P , say point Q .¹² Denote the locus of these points for different people by L , which we shall use as a mere line of reference. (Notice, however, that this line gives the full Pareto optimum associated to the prevailing allocation of tax burdens, if these were fixed as lump-sum transfers $\tau(w)$ and each consumer were allowed to move along 'his' 45° line, without facing any distortion.) It is easy to see that, as long as $\theta' \leq 1$, L lies northeast of θ and that if there is a point at which $\theta' = 1$ (i.e. zero marginal tax), θ and L must there coincide. If $\theta' < 1$ at the top, there must be a gap between θ and L at the top, given by the distance between the points a and a^* .

Our new net income schedule $\hat{\theta}$, improving the prevailing one θ , is now constructed as follows: we select a point such as \hat{a} between a and a^* and we induce the top man to 'stay' at \hat{a} by giving $\hat{\theta}$ the slope of *his* indifference curve through that point. We follow this direction towards the left, at each point giving $\hat{\theta}$ the slope of the indifference curve through that point of the man who, under θ , consumes *on the same 45° line*. Hence each and every consumer moves northeast along 45° lines and therefore $\hat{\theta}$ is feasible from the point of view of revenue (or of total output); by construction of $\hat{\theta}$ the allocation it implies is indeed chosen by consumers so that it is feasible from the point of view of de-

¹¹We are assuming, as was implicitly required also in the tax change of fig. 1, constant returns to scale, so that wages do not change with the labour supply. The assumption is removed below.

¹²If consumers' preferences were such that zero leisure could be observed at certain (positive) prices, the point Q could correspond to a corner-solution, which I exclude as an uninteresting (but easy to deal with) complication.

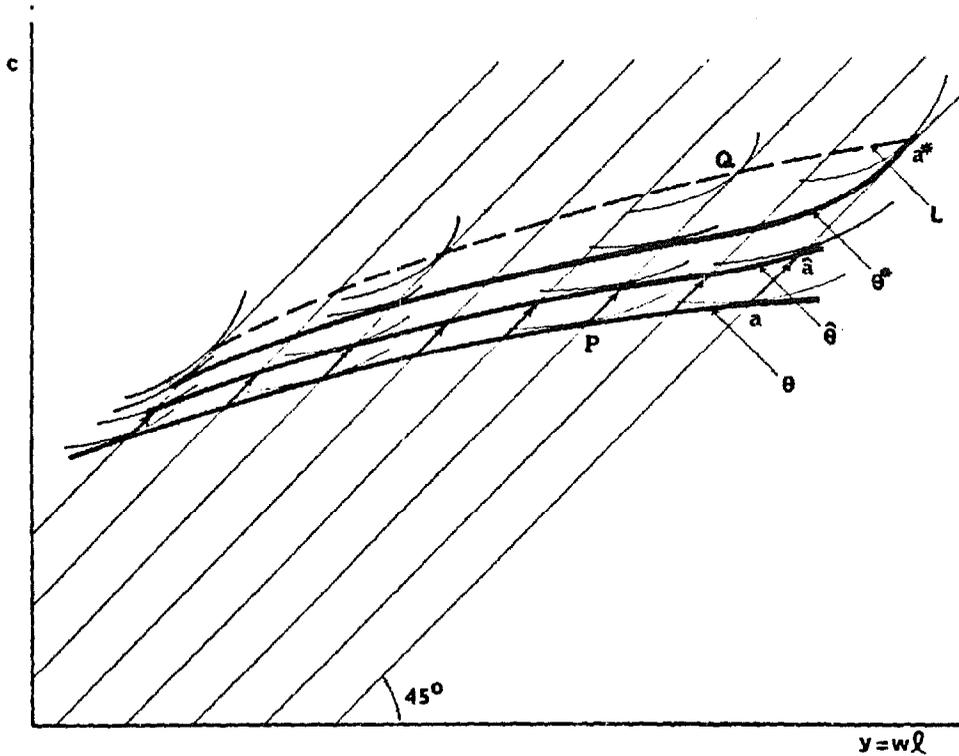


Fig. 2. A better ($\hat{\theta}$) and the best (θ^*) net income schedules preserving the ruling allocation of tax burdens.

centralization of consumption decisions too and, finally, $\hat{\theta}$ is clearly better than θ for each consumer. Hence the theorem.

A point that needs to be checked to validate the above argument is whether $\hat{\theta}$ must in fact lie everywhere strictly above θ and does not cross it or merge with it. That $\hat{\theta}$ cannot cross θ is immediate, because if there were a point of contact between the two curves, $\hat{\theta}$ would by construction coincide with θ from that point downwards. Hence the change is Paretian, no one is made worse off. On the other hand, it is shown in the appendix that as long as marginal taxes are strictly positive at every point of θ , $\hat{\theta}$ lies strictly above θ , so that it in fact leaves everybody better off. However, if at some point θ has a zero marginal tax, it is obvious that θ , L and therefore also $\hat{\theta}$ must merge together at the highest such point, $\hat{\theta}$ and θ obviously coinciding from that point downwards. Such a point would be a singularity for the family of variations we are considering, and this imposes the main part of the qualification of the scope of the theorem written in brackets in its statement. Similarly, if under θ there are a number of people not doing any work, then we must build $\hat{\theta}$ with a zero-marginal-tax range along the 45° line passing through the bottom ($y = 0$) of θ , so that anyone of the previously idle who now (under $\hat{\theta}$) decides to work pays, just as everyone else, the same amount of tax as before. But some of this group may still choose

not to work under $\hat{\theta}$, thus ending up in the same position they were under θ . Hence the restriction to positive incomes in the theorem.

The assumption of constant returns to scale has thus far been needed so as to leave the level and distribution of wages fixed while changing the aggregate supply of labour. It is easy to drop this assumption without causing any harm to the results – only to the simplicity in computing a given $\hat{\theta}$. For this purpose, we reinterpret w as being the *nominal* wage rate (relative to the wage rate of a reference man), which is given exogenously by the distribution of skills. Real wages are now ωw , where ω is the basic real wage rate (per efficiency unit); $y \equiv wl$ is now nominal income. Indifference curves between (real) consumption and *nominal* income can still be drawn in the diagram and are invariant to changes in ω . The only change that would need to be introduced in fig. 2 is that the technical transformation of changes in nominal income into equal changes in consumption is now not performed along straight lines of fixed, unitary slope. Instead, we first consider infinitesimal shifts of θ ,¹³ for which the slope of the transformation lines is the wage rate prevailing then, under θ . Call it $\omega(\theta)$. For larger (finite) changes, we have that the further our choice of $\hat{\theta}$ is from θ , the more the associated basic wage rate $\omega(\hat{\theta})$ falls (if decreasing returns prevail). Therefore, the lines that transform changes in efficiency units into equal changes in output (i.e. nominal income into consumption) are now not straight lines but certain concave curves, all having a slope $\omega(\hat{\theta}_v)$ at the points they cross the member $\hat{\theta}_v$ of the family of curves we have defined. Now, for purposes of computation, to find a particular curve $\hat{\theta}$ that represents a finite improvement on the initial net income schedule, all that is involved is the calculation of a succession of such $\hat{\theta}_v$'s along these concave curves with slope $\omega(\hat{\theta}_v)$. At each step we know the prevailing basic wage rate $\omega(\hat{\theta}_v)$, which was found in the previous step (starting the series with the wage rate actually ruling under θ), and which we use to compute $\hat{\theta}_{v+1}$ and the associated aggregate supply of labour, which yields $\omega(\hat{\theta}_{v+1})$, an input for the next computation.¹⁴ It is in any case unlikely that as a result of the tax changes we are describing the labour supply should change enough to introduce serious nonlinearities in the system – perhaps the assumption of constant returns or an approximation with a small number of linear steps will be reasonably accurate, all the more so from a long-run point of view, where the stock of capital is allowed to adjust to

¹³It is easy to see that no two $\hat{\theta}$'s can cross each other and that this implies that we can find one everywhere arbitrarily close to θ .

¹⁴In the foregoing, I have restricted attention to the direct interrelation between tax changes, changes in the supply of labour and the wage rate. However, when the supply of labour increases and there are diminishing returns, there is a corresponding increase in the level of profits, which means (i) that profits-recipients will benefit more than what our analysis would indicate, and (ii) that they will therefore revise (probably reduce) their supply of labour. Of course all this will not happen if all extra profits are taxed away and destroyed, say, so that our picture gives a valid distribution of *minimum* benefits. I thank D. Soskice for calling my attention to the presence of profits in this case.

changes in the labour supply. In fact, Flemming (1976) shows that, under certain conditions, purely redistributive earned income taxes do not affect the steady-state basic wage rate at all.

The tax schedules so constructed (whether with a constant or with a falling basic wage rate), form a one-parameter family of nonintersecting curves, each one of them Pareto-superior to all those below it. The parameter that defines them is simply the choice of the starting point \hat{a} , between z and a^* . On each of these curves each consumer pays the same total amount of taxes as under θ , while earning higher incomes. The incentive for this expansion is provided by a suitably selected general decrease in marginal tax rates¹⁵ while keeping the basic poll transfer unchanged. The necessary and sufficient condition for us to be able to construct the exercise is that there initially be a positive marginal tax rate at the top.

Clearly, from among the curves in this family, the best one is that which starts precisely at a^* , i.e. the one with a zero marginal tax at the top.¹⁶ Let us denote it by θ^* . Presumably the government chose to impose the currently ruling tax schedule simply because it is one that it likes for some reason. If so, and if the government is Paretian, then there is something to be said for θ^* , a tax reform that leaves the structure of tax payments essentially unchanged while allowing a strong Pareto improvement to take place. On the other hand, the government might in some cases be reluctant to impose a true (utilitarian) optimal tax schedule even if in principle it believed in it, particularly if it were very different from the taxes ruling at the time. For example, institutional (or other) barriers may still impede the implementation of a negative income tax. The second best may not be practicable but maybe 'the third' is, θ^* , which should not be so difficult to introduce as it works to everyone's advantage. Finally, a convenient feature of this tax reform is that, for purposes of actual computation, we do not need to investigate the underlying distribution of skills – perhaps the ingredient of optimal income taxes which is the hardest one to obtain – nor do we require specification of a welfare function (one is 'sanctifying' the current choices of the government and simply improving on them). Only knowledge of the work-behaviour of consumers (or of 'typical' consumers) remains necessary – and the tax schedule one is to start from. An appendix gives the differential equation form of this tax reform.

3. The bottom end of income taxes: Should some people be induced not to work?

It is tempting to apply similar arguments to the analysis of the bottom end of

¹⁵It is in fact a general *decrease* in the marginal tax rates faced by different people if, as we normally expect, indifference curves for each w become *steeper* as we move northeast along the relevant 45° line, in going from θ to $\hat{\theta}$. A sufficient condition for this to hold (below the line L) is noninferiority of leisure.

¹⁶This curve also maximizes aggregate output relative to the family defined by θ . Note that points a^* in figs. 1 and 2 are but the same.

the scale, particularly because it may in a relevant sense be 'more important' to have clear ideas about how we want to treat people in low income brackets. However, one soon realizes that the type of direct arguments used so far does not work at the bottom of the scale. Intuitively, the difficulty lies in that, other things being equal, reductions in marginal tax rates from any point *downwards* in the scale represent *increases* in tax payments by the people affected (from B^0 to B^1 in fig. 3) and is therefore not a good change for them, whereas reductions in marginal tax rates from any point upwards represent reductions in tax payments by those affected (from T^0 to T^1) and hence is an unambiguous improvement if feasible.

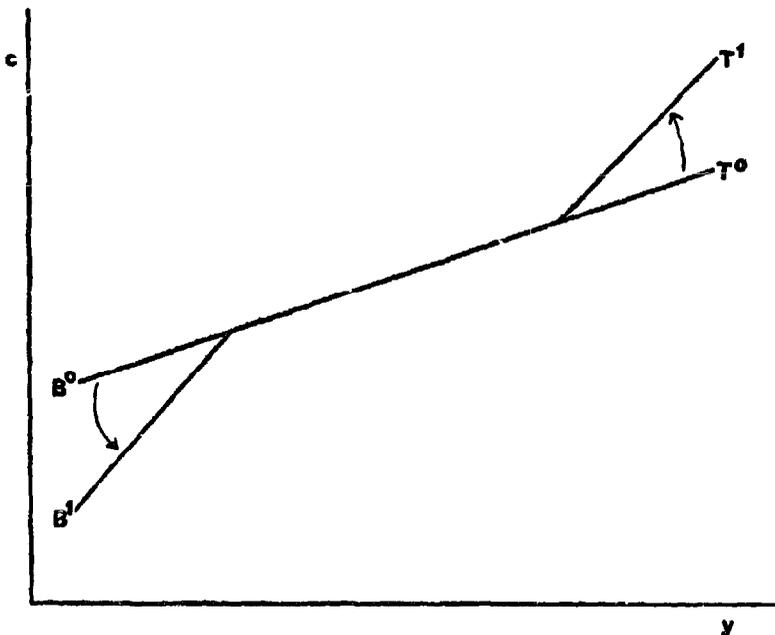


Fig. 3. Opposite effects on tax levels of reductions in marginal taxes.

When unable to use direct arguments to prove optimality or suboptimality of a candidate tax function, one must resort to the analysis of necessary conditions for an optimum, on the assumption that it exists. It will be shown in the following sections that this approach is particularly fruitful in the present context, for it yields a fairly general result (on the marginal taxes required at the endpoints of any nonlinear tax schedule) which in particular says that the property of the top of income taxes discussed in the last section applies to the bottom as well: if no household earns zero income in the optimal arrangement, then the optimal marginal tax faced by the lowest earner is zero. The condition at the top was boundedness of income; the condition at the bottom is the symmetrical one: boundedness away from zero (in the optimum). In fact, it will be

suggested in section 6 that in many cases all one needs for the result to hold at the bottom of the scale is that income be positive for all consumers whose wage rate is positive – i.e. we can often allow for the distribution of income to start off from zero, but we exclude the possibility that there be a number of people with different wage rates doing no work in the optimum. This requirement will probably be surprising to all those familiar with the optimal income tax literature and with the well-known result on the desirability of having ‘bunching’ of able people remaining idle at the bottom of the optimal income scale, a result which we now turn to discuss.

Consider an indifference curve between consumption and hours worked, which on the assumption of identical preferences is the same for everybody in the population, panel (a) of fig. 4, and consider the resulting map in consumption–income space shown in panel (b), obtained by multiplying each value of l in panel (a) by w . That is, in going from (a) to (b), one is stretching or compress-

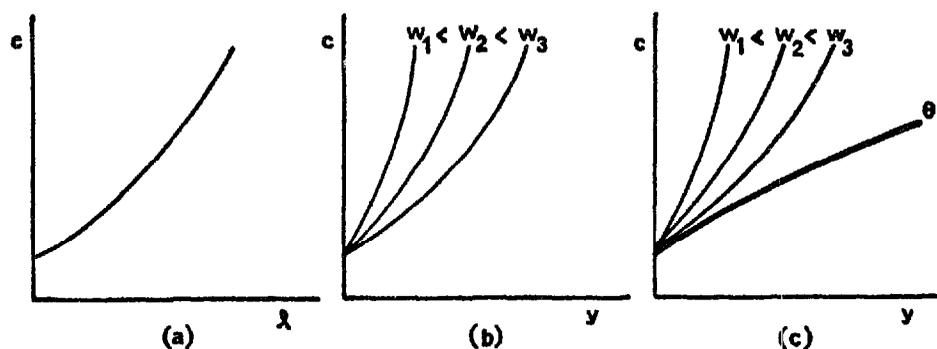


Fig. 4. ‘Bunching’.

ing horizontally all indifference curves by a factor w . It is immediate that if the marginal rate of substitution of c for l at $l = 0$ is strictly positive, as shown in fig. 4a, the resulting indifference curves in panel (b) must have a slope at each y (including $y = 0$) that increases without bound as $w \rightarrow 0$. The result of this is that when *any* net income function is drawn on the diagram, there always is a positive wage rate below which everybody chooses not to work, as long as $\theta'(y) \leq 1$ for all y , a prerequisite for optimality. This is shown in fig. 4c. Therefore, with preferences being as in fig. 4a, ‘bunching’ at the bottom *must* arise whenever the distribution of wages goes down to zero (and is continuous, or ‘nearly’ so, near the bottom), provided we restrict attention to nonnegative marginal tax functions.

The above result was noticed by Mirrlees (1971, p. 185) who, however, was implicitly making an assumption (see footnote 19, below) on individual preferences: that even when no work is performed ($l = 0$), individuals have a strict marginal dislike for work, and so behave as if they would like to work

even less, so to speak (the standard corner solution) – which is the case portrayed in panel (a) of fig. 4. But the diametrically opposite result can be arrived at by making what might appear to be a minor change in our description of the work-behaviour of ‘typical’ consumers. We may assume that people indeed always prefer having more rather than less leisure time, but that they are *satisfied* on that count (satiated of leisure) when they do no amount of work whatever, by which I mean that their marginal rate of substitution of c for l is zero at $l = 0$. Indifference curves would then meet the c -axis at right angles, as in panel (a) of fig. 5.¹⁷ But then, when curves like this one are stretched or compressed to obtain the maps on (c, y) -space of fig. 5b, it is clear that they retain their zero slope at $y = 0$. It then follows that any tax function will induce everybody (except people with $w = 0$)¹⁸ to do some work, provided only that the marginal tax rate at the bottom does not reach 100% ($\theta'(0) = 0$). We can notice in passing that although a net income function with $\theta'(0) = 0$ would probably be subopti-

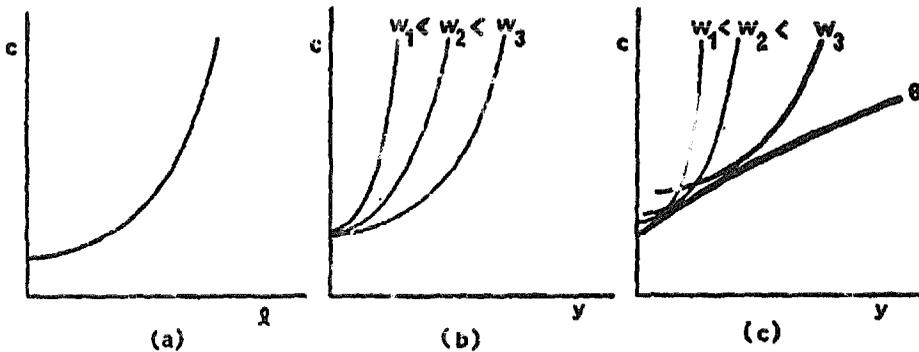


Fig. 5. No ‘bunching’.

mal in most interesting cases (for the extreme disincentive to work it imposes on the bottom fractions of the labour force), it nevertheless is precisely a feature of tax schedules commonly observed in practice (with unemployment benefits which are lost if any income is earned), so that it is just natural to expect some bunching at the bottom – induced unemployment – in actual economies, regardless of the type of work preferences that prevail.

We can then distinguish two main cases:¹⁹

¹⁷Two illustrative examples of utility functions having these indifference curves are given in section 6 (eq. (21)).

¹⁸The diagram only says that $y(w) \equiv wl(w) = 0$ at $w = 0$, but not that in fact $l(0) = 0$. However, this is obvious, because a 0-man’s budget constraint in (c, l) -space (i.e. in fig. 5a) is a horizontal straight line.

¹⁹It may be convenient to put briefly the foregoing in symbols. Let the utility function be $u(c, l) = u(c, y/w)$. Define the marginal rate of substitution of c for l (panels (a)) as $-u_l/u_c$ (with a minus sign to make it normally positive) and that of c for y (panels (b)) as $-u_l/wu_c$.

Proposition 1. (i) 'Bunching': if the marginal rate of substitution of consumption c for time worked l remains strictly positive as $l \rightarrow 0$, then there always is a certain number $w^0 > 0$ such that all consumers with $w \in [0, w^0]$ do not work, provided only that they face nonnegative marginal taxes (which includes the optimum). (ii) No 'bunching': if the marginal rate of substitution of c for l tends to zero as $l \rightarrow 0$ and if marginal taxes are less than 100% (at least at the bottom of the scale), then the only people who do not work are those who would not earn any income if they worked.

In either case it is also clear from the figures that if it is optimal for a w^0 -man not to work, this is also optimal for all $w \leq w^0$ [Mirrlees (1971, proposition 1)], although under (ii) this set consists only of $w^0 = 0$.

Which of these two basic types of work preferences is a more reasonable simplification of reality? The question is simply whether work-behaviour is such that individuals prefer not to earn any income by their own effort if the wage rate is low enough (but positive), choosing merely to rely on the basic poll-subsidy paid by the government. An important notion here is 'leisure', and what we take it to represent, i.e. whether it corresponds to the time not spent in production activities, which is perhaps its more natural definition, or whether it is the time not actually sold in the market as a service, which is the relevant notion for the fiscal authority interested in taxable incomes. Because in the second case, leisure includes possible nonmarket production activities, and it is clear that there may be some individuals whose productivity in these activities is greater than what they would earn as workers in the market, so that they withdraw their supply of work and become nonworkers in the eyes of the government. Of course these people will normally be wanting to sell some of their 'home' produce, so as to be able to buy other goods, in which case they are not zero-income bunchers in the end. Bunching will arise on account of the present argument only if non-market production is possible and if, in the face of a given tax schedule, people choosing not to sell their labour in the market also prefer to keep their home product rather than selling part of it. Needless to say, this story does not fit the simplified economy modelled above, for we are here referring to different *types* of production activities, in some of which incomes are not taxable at the production point (but only if sales are realized). Our interest is simply to seek a rationalization for or against the possibility of zero-income bunching on these lines.

Write $u^0 \equiv u(c^0, 0)$. If $-u_l^0/u_c^0 > 0$ (for $c^0 > 0$, $c^0 \equiv$ basic poll subsidy) as in fig. 4a, then as $w \rightarrow 0$, $-u_l^0/wu_c^0 \rightarrow \infty > \theta'(0)$ (because the latter is bound to be ≤ 1), and it pays all consumers from some w^0 downwards not to work. This is the case implicit in Mirrlees (1971) when he says (at the foot of p. 185) that unless $c^0 = 0$, u_l^0/wu_c^0 is unbounded as $w \rightarrow 0$. Our second case is when $-u_l^0/u_c^0 = 0$, as in fig. 5a, in which case $-u_l^0/wu_c^0 = 0$ for any $w > 0$, so that if $\theta'(0) > 0$, it pays all consumers with $w > 0$ to work.

At another level, we can assume away the existence of any nonmarket production activities, so that leisure consists only of nonproductive 'free time' (i.e. time not engaged in the production of marketable goods). It may reasonably be argued that more leisure time not only provides one with more rest away from work, say, but that it actually enables one to enjoy more one's consumption, which hints at indifference curves of the first type, as in fig. 4. This argument seems to me persuasive if the level of basic consumption one has in mind is high (i.e. when one is referring to indifference curves that meet the c -axis at high c^0), which is probably not the relevant case in the present context, where c^0 is just the poll subsidy, which cannot be particularly high at any rate. In our case it seems more likely that the consumer will not have much interest in increasing his leisure beyond a certain point, but will always try to sell some of his services in the market rather than being entirely idle all of his time.²⁰ One can pose the question: if there was a perfect labour market and consumers faced no income taxes but only a moderate fixed poll grant somehow financed, say, would there be voluntary unemployment? Only a negligible amount, I believe – of course leaving aside nonparticipation of a proportion of women (as well as of children and the aged), largely due to noneconomic factors.

4. Optimal nonlinear taxes

We have so far restricted attention to the taxation of (earned) income. More generally, however, we want to select an optimal tax mix: if say wealth, expenditure or perhaps food taxes are to be introduced, it is clear that their tax schedules should not optimally be selected in isolation from the income tax, but as the outcome of a single decision process, recognizing that one tax schedule's redistributive force and its effects on aggregate efficiency compound with those of the other tax schedule(s) present. The model of this section admits these possibilities, i.e. it can be interpreted as dealing with any one of the above-mentioned taxes or with a mix of them.²¹ Our main aim here is to derive an interesting qualitative statement on the shape of optimal nonlinear tax schedules in this fairly general setting.

The analysis of the previous section illustrated the delicate relation between assumptions on preferences and the qualitative nature of the optimum. Indeed, by merely giving a positive tilt to the bottom end of indifference curves, we went from having all consumers conveniently distributed along the tax curve, to a

²⁰Moreover, there is no reason why the level of leisure satiation should not occur at a positive level of work done, in which case (ignored above to avoid overcrowding of diagrams and possibilities) we would of course have again 'no bunching', with the curves of fig. 5 being a particular case.

²¹The model is the particular case $H = 1$ of a more general formulation presented in Seade (1977), where individuals differ in H dimensions. Mirrlees has also studied the more general problem; see, e.g., (1976).

situation in which a number of them are necessarily not working, 'bunching' at the bottom of the schedule. It is not difficult to see how this lesson extends to either end of other types of taxes. All we need is to ask whether there is a high enough price in the case of a consumption good (low enough in the case of work) as to choke off demand completely. The answer is no if each marginal rate of substitution of the numeraire for a commodity (or service) tends to infinity (zero) as the consumption of that commodity (or supply of work) tends to zero. These conditions are, in this sense, equivalent.

The relevance of the above discussion lies in that the result we shall derive below on the shape of optimal taxes requires, basically, that there should be no bunching of a nonzero proportion of people (if these are distributed continuously) at an endpoint of the corresponding tax schedule, as would result if their different maximization problems had a corner solution, so that we restrict attention to interior optima for consumers (although of course we can allow for corner solutions to arise at some prices, as long as in the optimum this is not the case). This would certainly be a very restrictive condition if the tax schedules put into the model were interpreted as referring to specific, disaggregated commodities, some of which would normally not be consumed at all by a considerable number of people. But in such cases the very setting up of a theory contemplating only nonlinear tax treatment of the consumption of these commodities is itself open to question, for no one I think would propose the introduction of nonlinear taxes based on the consumption of individual goods, whose trading among consumers would be impossible (or very expensive) to avoid. What I have in mind is to interpret our model and results mainly as applying to major aggregates which are usually or can conceivably be taxed nonlinearly, such as expenditure at different periods, income or wealth, the arguments in utility functions being such (e.g. consumption in certain years) that no corner solutions for individuals can sensibly exist.

4.1. *The model*

We shall assume that consumers can be fully described by a single parameter h (the wage rate, an index of needs, preferences of some sort) which follows a distribution function with density $f(h)$, defined for values of h in a certain interval:

$$h \in [\underline{h}, \bar{h}].^{22}$$

²²I take this interval to be the support of $f(\cdot)$, i.e. the smallest (closed) interval containing all the values of h for which $f(h) > 0$. When h is interpreted as an ability index, it is most natural to take it to represent the (relative) wage rate, so that \underline{h} and \bar{h} are fixed and given, but more generally these values are largely arbitrary and any strictly increasing transform of the index h (onto $[0,1]$ or $[0, \infty]$, say) could be used.

Since by assumption consumers differ only in h , they must otherwise have the same preferences, which we assume can be captured by a utility function, a particular cardinalization of which is selected by the government so as to reflect its own (individualistic) preferences:

$$U^h = u(\mathbf{x}(h), h), \quad (1)$$

where $\mathbf{x} = \mathbf{x}(h)$ is the vector of N commodities consumed by h (although later we shall distinguish a numeraire x_1 and still denote by \mathbf{x} the remaining commodities). I assume that, for each h , $u(\cdot, h)$ is a continuous, monotonic, strictly concave and everywhere differentiable function of $\mathbf{x}(h)$ and that preferences (both utilities and marginal utilities) vary continuously with h .

Let us formulate the tax problem in a general way: the government imposes a (closed) set Θ of possible after-tax consumption bundles on consumers, from which the latter must make a choice:

$$\mathbf{x}(h) \in \Theta, \quad \text{Choice Set.} \quad (2)$$

The relevant part of the frontier of Θ (i.e. its consumption-efficient points) we denote by θ , which we take to be piecewise differentiable.

The government is interested, we assume, in maximizing the sum of individual utilities:

$$W \equiv \int_{\bar{h}}^{\underline{h}} u(\mathbf{x}(h), h) f(h) dh, \quad \text{Social Welfare.} \quad (3)$$

There are two complications that need not be introduced explicitly: the first is that in fact the government may be interested in raising or disposing of a certain revenue. We assume that the welfare value of this revenue (for the expenditure it may finance, say) enters a more general social welfare function in a weakly separable fashion, so that the government is still interested in maximizing (3) over the possible allocations $\mathbf{x}(h)$. The second is that, seemingly, we should allow for any strictly quasi-concave transformation of utilities $G(U^h)$, or more generally $G(U^h, h)$, to appear in the maximand (3), so as to reflect different possible attitudes towards inequality on the part of the government. However, this is an uninteresting way of complicating our notation, for the particular cardinalization $u(\cdot, \cdot)$ was itself picked by the government so as to reflect exactly its judgement (while respecting individual ordinal preferences).

The maximization of (3) is limited by some constraints: first is the requirement that the aggregate demand for each commodity (including demands by the government, which I neglect) and supply of each factor be compatible with each other, given the technology available. This can be captured by the following twin constraint:

$$\int_{\bar{h}}^{\underline{h}} \mathbf{x}(h) f(h) dh \leq \mathbf{X}, \quad \text{Aggregate Demands;} \quad (4)$$

$$\varphi(\mathbf{X}) \leq 0, \quad \text{Production Possibilities.} \quad (5)$$

Secondly, the government cannot directly choose the allocations $x(h)$, but has instead to rely on the working of a tax scheme. Hence, it must take into account the choices $x(h)$ individuals will make, as utility maximizers, in the face of the taxes. This is what makes the problem a second-best one. To give a manageable form to this condition imposed by the need to decentralize decisions notice that an optimal consumption plan for an h' -man must maximize his utility over the consumption plans chosen by people with any other values of h , so that the constrained problem

$$\max_x u(x, h') \quad \text{s.t.} \quad x \in \theta,$$

can be expressed as the unconstrained one

$$\max_h u(x(h), h').$$

That is, an h' -man can in principle *behave* as an h' -man does (in what concerns his tradings with the market, so that it is not hours worked that count here, but income actually earned, for example), as long as he also pays the same taxes paid by h' . By definition, the best policy for h' to follow is to behave as he finally does, as h' himself, so that the last expression reaches a maximum at $h = h'$. This maximization, in the absence of corner-solutions by individuals and of corners of $x(h)$,²³ requires that

$$u_x \cdot x' = 0. \quad \text{Individual Behaviour,} \tag{6}$$

where x' is the vector of derivatives of $x(\cdot)$ at h , while u_x are consumption-partial derivatives of $u(\cdot, \cdot)$. What this equation ultimately imposes is equality of each marginal rate of substitution of an h -man with the corresponding partial derivative of θ at the point he chooses, that is, that θ should be an envelope of the relevant indifference curves at the relevant points.

A third type of constraint that should be introduced (and which would yield ad hoc necessary conditions holding at points where it becomes binding) is non-negativity of $x(h)$. However, the qualitative result I wish to concentrate on below applies to cases where this constraint is not binding and we can therefore ignore it here. The same applies to second order conditions for individual maximization, which are assumed to be satisfied.

²³I avoid the complication of explicitly allowing for corners of the allocation functions contained in $x(h)$ because it will be enough for the result in section 5, our main interest here, to have their description where they are smooth, for the reason that a piecewise smooth function must be smooth over a sufficiently small half-interval near any point. This is all we need to be able to take limits as $h \rightarrow \underline{h}, \bar{h}$.

4.2. The maximum

We wish to find a maximum of (3) subject to the constraints imposed by production limitations and by the need to decentralize decisions, (4), (5) and (6). We can construct a Lagrangean by defining a scalar multiplier λ^0 for (5), N scalar multipliers λ for (4) and a function-multiplier $\mu(h)$ for (6), as this constraint must hold at each h . Hence our problem is equivalent to the unconstrained maximization of:

$$\mathcal{L} \equiv \int_{\underline{h}}^{\bar{h}} F(h, \mathbf{x}, \mathbf{x}') dh + [\lambda \cdot \mathbf{X} - \lambda^0 \varphi(\mathbf{X})], \quad (7)$$

where

$$F \equiv F(h, \mathbf{x}, \mathbf{x}') \equiv \{[u(\mathbf{x}, h) - \lambda \cdot \mathbf{x}]f + \mu u_{\mathbf{x}} \cdot \mathbf{x}'\},$$

by choice of $\mathbf{x}(h)$, \mathbf{X} and the multipliers.

An obvious necessary condition is, as one would expect, that

$$\lambda = \lambda^0 \varphi_{\mathbf{x}}(\mathbf{X}), \quad (8)$$

that is, that there should be production efficiency, sustained by equating relative shadow (or producers') prices to marginal rates of transformation. That in what concerns the aggregate variables \mathbf{X} . Let us now turn to the disaggregated variables, $\mathbf{x}(h)$. Euler's necessary condition for a maximum of (7) is that

$$F_{x_i} = \frac{d}{dh} F_{x'_i}, \quad (9)$$

and since

$$F_{x_i} = (u_{x_i} - \lambda_i)f + \mu u_{x x_i} \cdot \mathbf{x}'$$

and

$$F_{x'_i} = \mu u_{x_i} \Rightarrow \frac{d}{dh} F_{x'_i} = \mu' u_{x_i} + \mu u_{x x_i} \cdot \mathbf{x}' + \mu u_{x_i h},$$

our Euler equation (written in vector form) is

$$(u_{\mathbf{x}} - \lambda)f = \mu' u_{\mathbf{x}} + \mu u_{\mathbf{x}h}. \quad (10)$$

This equation (together with individual behaviour (6) and information on the endpoints, to be discussed presently) in principle describes the required pattern

of distortions, the taxes. However, if we were to double all prices, when nothing should in principle change, the values in the left-hand side of (10) would in fact change, showing an undesirable dependence of the outcome on the particular choice of the price level. We need a price-invariant version of these equations, where prices appear as relative prices and marginal utilities as rates of substitution. For this purpose, let us introduce a slight change in our notation, so as to distinguish a numeraire, the first commodity, from the remaining ones, which are now denoted by \mathbf{x} themselves. Hence, the vector-equation (10) now represents the Euler equations for commodities, 2, . . . , N , while for the numeraire we have

$$(u_1 - \lambda_1)f = \mu'u_1 + \mu u_{1h}, \tag{11}$$

where u_1, u_{1h} denote u_{x_1}, u_{x_1h} (I shall also write u_i, u_{ih} for u_{x_i}, u_{x_ih} , below).

Denote by $\mathbf{s} \equiv \mathbf{s}(x_1, \mathbf{x}, h)$ an h -man's vector of marginal rates of substitution of the numeraire for each commodity, that is:

$$\mathbf{s} \equiv \frac{u_{\mathbf{x}}}{u_1}, \tag{12}$$

whose partial derivatives with respect to h are

$$\mathbf{s}_h = \frac{u_{\mathbf{x}h}}{u_1} - \frac{u_{1h}}{u_1} \mathbf{s}. \tag{13}$$

Using (12) and (13), eqs. (10) become

$$(u_1 \mathbf{s} - \lambda) f = \mu' u_1 \mathbf{s} + \mu (u_{1h} \mathbf{s} + u_1 \mathbf{s}_h),$$

and subtracting (11) multiplied by \mathbf{s} from this last expression we get

$$\left(\mathbf{s} - \frac{\lambda}{\lambda_1} \right) = \frac{u_1 \mathbf{s}_h}{\lambda_1 f(h)} \mu(h). \tag{14}$$

The left-hand side of eq. (14) gives the distortions that should optimally be imposed on an h -man's choices between each commodity and the numeraire, for it shows the differences between the relative prices he faces (to which he equates his marginal rates of substitution \mathbf{s}) and the prices he would face in the absence of intervention by the government. In fact, $(\mathbf{s} - \lambda/\lambda_1)$ is precisely the vector of *marginal taxes* levied on the consumption of each commodity in terms of numeraire in the optimum.

We can notice in passing that if the utility function has additive separability of the form

$$u^1(x^1, h) + u^2(x^2),$$

where $(x^1, x^2) = x$, then an h -man should pay the same marginal taxes on his consumption of each commodity in the group x^2 , i.e. no distortion of his choices should be introduced within that group (for each h). This can be seen taking as our numeraire any one commodity in x^2 , in which case s has no explicit dependence on h (and hence $s_h = 0$) for all the other commodities in that group, and the distortions in (14) are identically zero.²⁴ In particular, if we interpret the model as referring to (single-period) income and excise taxes and if there is separability between labour and all other commodities, then no commodity taxation should be imposed at all. Notice that, since this is true in the present case where excise taxes are allowed to take any shape, it is of necessity true when they are restricted to proportionality. Notice too that the result still holds in an intertemporal context if labour in each period is separable from the other commodities (and abilities remain unchanged with age in relative terms), so that under these assumptions there should be no interference with intertemporal decisions either.

We still need to get rid of the multiplier $\mu(h)$ in (14). For this purpose we notice that we still have one Euler equation available, say (11) (the $N-1$ equations (10) were used up in deriving the $N-1$ we have in (14), so to speak), which ultimately is a linear differential equation for $\mu(h)$. Solving it:

$$\mu(h) = \int_h^{\bar{h}} \left(1 - \frac{\lambda_1}{u_1}\right) \exp\left(-\int_h^{h'} \frac{u_{1h}}{u_1} dh''\right) f(h') dh' + \mu(\bar{h}). \quad (15)$$

The value of the constant $\mu(\bar{h})$, as well as the value of $\mu(h)$, are shown below to be zero in most cases of interest.

Eqs. (6), (14) and (15) form a system of $N+1$ differential (and integral) equations in the $N+1$ unknown functions $x_1(h)$, $x(h)$ and $\mu(h)$. This provides a sufficient description of the optimum once we know which particular solution to select, which amounts to determining the values of x_1 , x and μ on either end of the distribution, at $h = \underline{h}, \bar{h}$. In calculus of variations problems such as the present one where the end-values of the independent variable are fixed exogenously ($h = \underline{h}, \bar{h}$ here), the standard transversality conditions take a very simple form: that F_x (F defined in (7)) should vanish at both ends [see, e.g., Gelfand and Fomin (1963, sec. 1.6)]. That is, we require that

$$\mu(h)u_x(x(h), h) = 0, \quad (16)$$

²⁴The result is due (in the form to be given presently) to Atkinson and Stiglitz (1974). Mirrlees (1976) generalized it by showing that still no distortion should be introduced within the group x^2 when only weak separability prevails, of the form $u(x^1, h, a(x^2))$.

and

$$(16')$$

$$\mu(\bar{h})u_x(\mathbf{x}(\bar{h}), \bar{h}) = 0,$$

both of which must hold for each commodity, including the numeraire. To conclude from this that $\mu(\underline{h}) = \mu(\bar{h}) = 0$, all we require is that *any* marginal utility be non-zero for a bottom (i.e. \underline{h}) man, similarly for a top man, that is, that neither of these two types of persons be fully satiated (i.e. with respect to every commodity and service) in the optimum. This clearly covers most cases of interest and we may safely conclude that

$$\mu(\underline{h}) = \mu(\bar{h}) = 0.^{25} \tag{17}$$

We lack a similar explicit result for the end-values of x_1, \mathbf{x} , because all the equations in (16) and (16'), the transversality conditions, become identities once (17) is introduced. Our information on how to specify the particular solution we are interested in in terms of x_1, \mathbf{x} comes from elsewhere: eqs. (5) and (8) (the production function and production efficiency) determine the aggregate total supply of each commodity. That is, rather than fixing end values for x_1, \mathbf{x} , we imagine the net-consumption function θ being placed at different levels for each commodity, until a particular solution of the system of differential equations (the micro, or allocation side) induces market clearance for all commodities, that is, yields aggregate demands (4) that solve eqs. (5) and (8), the supply side.

5. Zero marginal taxes at the endpoints of optimal schedules

Let us collect eqs. (12)–(17) for ready reference:

$$\left(s^i - \frac{\lambda_i}{\lambda_1}\right) = \frac{u_1 s_h^i}{\lambda_1 f(h)} \int_{\underline{h}}^h \left(1 - \frac{\lambda_1}{u_1}\right) \exp\left(-\int_{h'}^h \frac{u_{1h'}}{u_1} dh''\right) f(h') dh'; \tag{18a}$$

$$u_1 \mu \rightarrow 0 \quad \text{as} \quad \begin{cases} h \rightarrow \underline{h}, \\ h \rightarrow \bar{h}, \end{cases} \tag{18b}$$

(where $\mu \equiv \mu(h)$ is the integral term in (18a)); and

$$s^i \equiv \frac{u_i}{u_1}; \quad s_h^i = \frac{u_{ih}}{u_1} - \frac{u_{1h}}{u_1} s^i. \tag{18c}$$

²⁵In the income tax case (when h is the nominal wage rate) this equation *always* holds [it corresponds to (20) or (26) in Mirrlees (1971)]. When the distribution of skills is unbounded, so that $\bar{h} = \infty$, satiation 'at the top' may in some cases occur (in the sense that the amount of work done may be tending to zero yet income and consumption be tending to infinite with the wage rate), but (17) can still be deduced as a limit.

((a) comes from (14), (15) and (17), and (b) is (16) and (16') for the numeraire.)

I make the following assumptions, parts of which were already imposed above:

Assumption 1. $u(\cdot, h)$ is continuous, monotonic, strictly concave and everywhere differentiable in x . Besides, u_{ih} is bounded for $0 \ll x \ll \infty$. Some of the conditions in the first part can be relaxed without harm. The second part refers to the way in which the parameter h enters the utility function, requiring that the rate of change (with h) of the marginal utilities associated to a fixed bundle of goods should be bounded. It is only marginally a stronger requirement than one of continuity of $u_i(x, \cdot)$ as a function of h , basically asking that similar people (h -wise) have similar preferences or abilities (as is the case, e.g. in the income tax formulation).

Assumption 2. The density $f(h)$ does not oscillate infinitely often as h changes. In fact less would be enough, that $f(h)$ be of bounded variation, but in either case the requirement seems not to impose any significant restriction.

Assumption 3. The welfare function(a) is utilitarian and its integral over a zero-measure set of people is zero. The second part excludes the maximin valuation, in particular. 'Utilitarian' is taken in an extended sense, allowing the government to impose transformations on (or to weight differently) different people's utilities.

Assumption 4. The allocation vector-function $x(h)$ and the resulting tax schedule θ (i.e. $\theta(x) = 0$) are continuous and smooth. Unlike the preceding assumptions these impose certain conditions on variables which are endogenous to the problem, but the conditions do not represent any interesting economic restrictions. In fact they could be substantially relaxed, asking only that they hold in a local way, $x(h)$ being continuous and smooth near enough \underline{h} and \bar{h} , and $\theta(x)$ being smooth at $x(\underline{h})$ and $x(\bar{h})$ if these are ≥ 0 .

We can now state our central result, which extends the familiar zero-marginal-tax property of the top end of (bounded) income taxes to both ends of any optimal nonlinear tax or tax mix satisfying certain boundedness conditions:

Theorem 2. Consider an optimal tax regime and assume that assumptions 1-4 hold for \underline{h} . Then if $0 \ll x(\underline{h}) \ll \infty$, all marginal taxes must be zero at \underline{h} . Similarly for \bar{h} .

Notice that this result is stated independently for the top and for the bottom of the optimal tax regime, so that unboundedness of some x 's at the top, for instance, does not preclude the result from holding at the bottom. However, it will often hold at both ends, and in such cases the following is immediate.

Corollary. If the conditions of theorem 2 are satisfied at both ends of an optimal tax schedule (or set of schedules), the latter (or each of the latter) cannot be progressive nor regressive throughout.

Before we proceed to prove the theorem, it will be convenient to see why it should hold and when it would not. Eq. (18a), whose left-hand side gives the marginal tax function for commodity i , can be written as

$$\lambda_1 m_i = s_h^i(u_1 \mu) / f, \quad (19)$$

where m_i is the marginal tax on good i with respect to the numeraire. When we let h tend to \underline{h} (or to \bar{h}) in this equation, the term $(u_1 \mu)$ tends to zero, and there is a prima facie case to expect m_i to tend to zero as well. However, there are some reasons why this may not happen, notably: (i) If in the optimum there is bunching at the bottom of the scale (similarly for the top), so that $x(h)$ is identical for all h in a certain range $[\underline{h}, h^0]$ (with $h^0 > \underline{h}$, i.e. different people receiving equal baskets of goods at an endpoint of the tax schedule) then, as we move down the scale of h 's, the bottom of the tax schedule is first met at h^0 , not at \underline{h} , and at h^0 the term $u_1 \mu$ does not vanish. That is, 'bunching' at either end of $x(h)$ definitely upsets the result, calling for nonzero endpoint marginal taxes. The main condition that excludes here the possibility of bunching (other than an implicit assumption that h is a relevant parameter, so that unless there is a specific external inducement to the contrary, people with different h 's will normally choose different x 's) is the assumption that $x \geq 0$, so that consumers are not pressed against their nonnegativity constraints and choose interior maxima.²⁶ (ii) In the absence of bunching (and with assumptions 1-4 holding), the right-hand side of (19) tends to zero as $(u_1 \mu) \rightarrow 0$ (and so does the marginal tax m_i) whenever the ratio (s_h^i / f) does not diverge dominating $(u_1 \mu)$. Provided s_h^i remains bounded, it is shown below that divergence of f^{-1} never dominates $(u_1 \mu)$, so that the theorem holds regardless of the (direct effect of the) shape of $f(\cdot)$. On the other hand, if x remains bounded and strictly positive over all h , we do have boundedness of s_h^i as required. But if we allow some components of $x(h)$ or of $x(\bar{h})$ to be zero or infinite (still excluding bunching), then s_h^i may be unbounded and, if it is, this may make marginal taxes tend to nonzero limits as $h \rightarrow \underline{h}, \bar{h}$. This is the reason why marginal taxes do not always tend to zero in the upper tail of the unbounded income cases studied by Mirrlees (1971), where consumption generally tends to infinity as we move towards the 'top' ($h \rightarrow \infty$).

Remark. In many cases the endpoint zero-marginal-tax result still holds even if some x_i 's are actually zero or infinite at either end (illustrations for the bottom

²⁶Another, less central kind of bunching at an end-point can arise if θ is kinked in such a way that a number of (different) consumers at that endpoint all choose the same point $x(h) \geq 0$ on the kink as an interior maximum. This is excluded by the last part of assumption 4.

end of income taxes are provided in the closing section), although exceptions may then arise. What we basically exclude, however, are situations where there is bunching at either end of the tax schedule, as a result of a number of people having corner solutions for their individual maximization problems.

We can now proceed to the proof of theorem 2.

Proof. It is easily seen that strict concavity and monotonicity of $u(\cdot, h)$ (in assumption 1) together imply that u_i (each i) can vanish, if at all, only if, for some j , $x_j = \infty$ or $x_j = 0$ (usually when x_i itself is infinite if it is a 'good,' or zero if it is a 'bad'), which are excluded from the range of application of the theorem. Hence we may assume that in fact $u_i \neq 0$, $u_1 \neq 0$ for all allocations that may actually arise. In the same way it follows that $|u_i| < \infty$, $|u_1| < \infty$. Hence, $s^i \equiv u_i/u_1$ has $0 < |s^i| < \infty$. Boundedness of u_{ih} (in assumption 1), of s^i and of u_1^{-1} for all relevant allocations together imply that s_h^i is also bounded. Given boundedness of s_h^i , the theorem follows immediately from (18a), using (18b), if at either end $f(\cdot) \neq 0$. Thus assume $f(\underline{h}) = f(\bar{h}) = 0$. If we now imposed certain conditions on the behaviour of the function $f(\cdot)$ near the endpoints,²⁷ the theorem would follow immediately from here, by (repeated) use of L'Hôpital's rule. But we can more generally proceed as follows. Let us concentrate on the bottom end.

Boundedness of u_{ih}/u_1 implies that as $h \rightarrow \underline{h}$ (and since $h \geq h' \geq \underline{h}$ in (18a), this means as $(h-h') \rightarrow 0$), the exponential in the integrand of (18a) tends to one: $\exp(\cdot) \rightarrow 1$, and it certainly remains bounded. Now, from continuity of $x(h)$ (near enough \underline{h} ; assumption 4) and of $u_1(x, h)$, it follows that u_1 is a continuous function of h . Hence, by the sign-preserving property of continuous functions, $(1 - \lambda_1/u_1)$ is either zero or single-signed in $(\underline{h}, \underline{h} + \varepsilon')$ for low enough ε' .

Now consider $f(h)$. Since it does not oscillate infinitely often (assumption 2), there must be a small enough interval $(\underline{h}, \underline{h} + \varepsilon'')$ in which $f(h)$ is monotonic. Further, since $f(h) \geq 0$, it must be the case that $f(h) > 0$ in the whole of the open interval $(\underline{h}, \underline{h} + \varepsilon'')$, because since it is monotonic in that interval, it can only be zero at a point if it is zero at all points below that one, which contradicts our notion of \underline{h} as the infimum of the support of $f(\cdot)$ (the theorem would then apply to the lowest relevant value of h).

Hence if we take $\varepsilon = \min(\varepsilon', \varepsilon'')$, we have that in $(\underline{h}, \underline{h} + \varepsilon)$, the expression $[(1 - \lambda_1/u_1) \exp(\cdot)]$ is single-signed and bounded and $f(\cdot)$ is positive and attains its maximum at $\underline{h} + \varepsilon$. Therefore the absolute value of the integral in eq. (18a) is not decreased if we replace the $f(h')$ appearing in it by $f(\underline{h} + \varepsilon)$. That is, from (18a) for $h = \underline{h} + \varepsilon$,

²⁷That $f(\cdot)$ has all derivatives up to order k in some open intervals near its endpoints and that its one-sided derivatives of order k at the endpoints do not vanish, for some $k \geq 1$.

$$\begin{aligned} \left| \left(s^i - \frac{\lambda_i}{\lambda_1} \right) \right| &= \left| \frac{u_1 s_h^i}{\lambda_1 f(\underline{h} + \varepsilon)} \int_{\underline{h}}^{\underline{h} + \varepsilon} \left(1 - \frac{\lambda_1}{u_1} \right) \exp \left(- \int_{h'}^{\underline{h} + \varepsilon} \frac{u_1 h}{u_1} dh'' \right) \right. \\ &\quad \left. \times f(h') dh' \right| \\ &\leq \left| \frac{u_1 s_h^i}{\lambda_1} \int_{\underline{h}}^{\underline{h} + \varepsilon} \left(1 - \frac{\lambda_1}{u_1} \right) \exp(\cdot) dh' \right| \\ &\rightarrow 0 \quad \text{as } \varepsilon \rightarrow 0, \end{aligned}$$

from boundedness of s_h^i , of u_1 , of u_1^{-1} and of $\exp(\cdot)$.

It follows that $s^i(\underline{h}) = \lambda_i/\lambda_1$, that is, consumers with $h = \underline{h}$ face the same relative prices as producers – no taxes at the margin.

Exactly the same proof applies at the top, now working in an interval $(\bar{h} - \varepsilon, \bar{h})$ in which both $f(h)$ is monotonic and $(1 - \lambda_1/u_1)$ is single-signed or zero.

6. Concluding remarks

(i) The result just derived, being independent of the interpretation one may give to its various variables (i.e. of what the problem at hand is) stresses a fundamental similarity between different optimal nonlinear taxation models one may wish to consider. Nonlinear taxes, in the present stage of organization of modern societies, can include those on income, on wealth, on expenditure or even on important consumption goods whose retrading amongst consumers can be prevented, such as energy, perhaps. In all these cases the optimal arrangement makes the highest and the lowest consumers in the corresponding scale face undistorted shadow prices at the margin, at any rate at an end of the scale where the conditions of the theorem are satisfied. I expect similar results to hold in contexts other than taxation, such as optimal payment schemes or pricing policies, or other problems where incentives are central to the analysis and nonlinear policies can be pursued.

There is no doubt that some interesting examples are left out by the requirement that no consumption be zero or infinite at the ends of the scale. In these cases the result does not generally apply, although it may still explain the qualitative shape of optimal taxes mentioned in the introduction (typically progressive near the bottom and regressive towards the top). On the other hand, it is clear that, in general, the consumption of some goods (or supply of bads) being zero at either end of an optimal arrangement does not impede optimal marginal taxes from being zero there: it only opens the door to exceptions. One can in such cases proceed taxonomically, specifying the structure of the problem at hand and, for this, seeking conditions for which the result still holds

and perhaps others for which it necessarily does not. For the income tax problem, for example, the marginal tax rate is in many cases still zero at the bottom, even if we have people with zero skill and hence zero income, as long as we do not make *able* people earn *no* income through our assumptions on preferences (of the 'bunching' type of fig. 4). A heuristic argument can be provided for the easy case when $f(0) > 0$, for then (18a) behaves like:

$$\left(s^i - \frac{\lambda_i}{\lambda_1}\right) \approx \left(\frac{u_1}{\lambda_1} - 1\right) w s_w^i, \quad (20)$$

where w has replaced h and the index i denotes income (i.e. work done in efficiency units, whose price per unit, λ_i , is the basic wage rate, ω of section 2), with consumption being taken as good 1, the numeraire. Then, since consumption by a bottom man is positive in the optimum, u_1 is bounded, and therefore the condition for the marginal tax to be zero at the bottom is that $w s_w^i \rightarrow 0$ as $w \rightarrow 0$. This condition can be checked for particular cases. For example, it holds when utilities are additive separable in consumption c and hours worked l and the subutility of the latter is of the form

$$v(l) = (1-l)^\gamma/\gamma + l \quad \text{or} \quad v(l) = -e^l + l, \quad (21)$$

which are just two cases that satisfy $v'(0) = 0$ and therefore also $u_l/u_c \rightarrow 0$ as $l \rightarrow 0$ (the 'no bunching' condition on preferences), while at the same time the indifference curves they yield have the normal shape otherwise.²⁸ Of course if $f(0) = 0$ the result may be upset, the outcome depending on the speed with which $f(w)$ tends to zero as $w \rightarrow 0$. However, writing w for the lowest wage rate in the economy, it is worth noting that the result holds when $w > 0$ (with any $f(w) \geq 0$) and that it seems normally also to hold when $w = 0$ if $f(0) > 0$. Thus the occurrence that $w = 0$ and $f(0) = 0$, which is more likely to upset the result, seems to be the exception, perhaps more a feature of models we use than of actual economies. By and large, it transpires that the shape of the distribution of wages does not impose important qualifications on our result as applied to the bottom of optimal income taxes. The main source of potential limitations to the result in given cases seems to me to be the nature of preferences near the bottom: whether they give rise or not to an amount of bunching. Because, as was pointed out above, if some people with different wage rates all select the same point at the bottom of the scale, the relevant value of w in determining the marginal tax at that point is some w^0 , the largest w amongst nonworkers, at which value the

²⁸To verify that in these cases $w s_w \rightarrow 0$ as $w \rightarrow 0$, we write $r(c, l) \equiv u_l/u_c = r(c, y/w)$, hence $s \equiv u_y/u_c = r/w$, which by differentiation (using $l = y/w$) yields the following formula, general for the income tax case: $w s_w = -s(1 + \eta_{rl})$, where $\eta_{rl} \equiv (l dr/dl)/r$. We then calculate this expression for the examples given and let $w \rightarrow 0$. In both cases we find that $\eta_{rl} \rightarrow 1$ and $s \rightarrow 0$ as $l \rightarrow 0$ for each w and hence as $w \rightarrow 0$ (using $w \rightarrow 0 \Rightarrow l \rightarrow 0$, see footnote 18).

integral in (18a) (and hence the marginal tax) does not vanish. Notice, however, that in such cases, other things being equal, the value of this integral will be smaller and the zero-marginal-tax result more nearly met the smaller the proportion of people who remain idle in the optimum.

(ii) It is tempting to look for an intuitive explanation of the result in its general form. For this purpose, let us refer to the famous trade-off between efficiency and equality. The former, we know from basic Welfare Economics, requires that all consumers face the same prices, which must be the same as for producers. If it were possible to have the correct set of lump-sum transfers, one could first achieve 'equality' (the desired distribution of property, given the existing distribution of wages) and then allow full efficiency to take place by following the equal prices rule. But in the absence of lump-sum transfers we must be prepared to part with some efficiency for the sake of enforcing net transfers from rich to poor through the tax system. Now, when we change a marginal tax at a given point, it is not the person previously choosing that point who is primarily affected but, *prima facie*, all those above that person, for whom the whole tax schedule is raised or lowered. Hence, the smaller the number of people there are above a person in the scale, the weaker the distributive (or revenue) implications of the marginal tax imposed on him and the lesser the case to introduce a distortion there. In the limit, i.e. at the top, only efficiency arguments prevail and a zero marginal tax is called for. But a little thought shows that this reasoning in fact applies to both ends of the schedule, because an increase (say) of the marginal tax at any given point is as much a means of raising the tax schedule from that point upwards, thus exacting some more revenue from all higher incomes, as it is a means to lower the tax schedule from that point downwards, thus effecting a certain transfer per head to all lower incomes (precisely financed, given a revenue constraint, from the increased tax payments from all incomes in the upper portion). This is illustrated in fig. 6 by the tax function τ^0 and its modified form τ^1 . But now, just as equality considerations vanish when the top is reached because there are no further incomes to be taxed, they also vanish at the bottom because there are no possible beneficiaries of the tax change further down the scale, so that a change of the bottom marginal tax rate affects nothing but efficiency (incentives), which should therefore be given all the weight.²⁹

This argument has for verbal convenience been put in terms of an income tax, but in fact it applies to the general case, for allocative efficiency in the choices of each person always requires shadow-price pricing, and if we want to impose any

²⁹If, however, we attach a strictly positive (i.e. of positive measure) welfare value to the utilities of people *at* the very endpoint of the tax schedule, as is the case if there is bunching at the bottom or, indeed, under maximin, then the above argument does not apply (the equality gain from the tax change does not vanish as we approach the bottom of the schedule) and a positive endpoint marginal tax is called for on distributional grounds, so as to 'lift' the level of basic consumption for any given shape of the rest of the schedule.

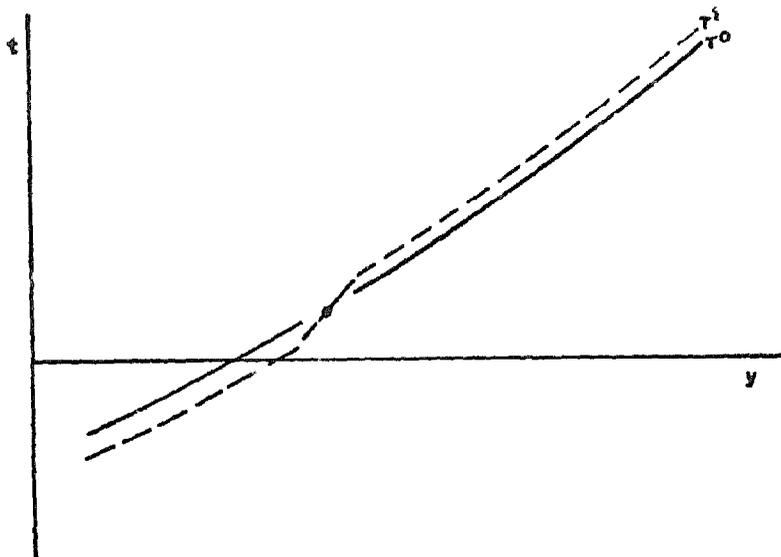


Fig. 6. A revenue-neutral marginal tax change at a point.

other set of prices on an h -man (i.e. nonzero marginal taxes in some of his tradings with the market), it must only be because we expect in that way to achieve some redistribution in a certain desired direction, such as from rich to poor (which is the usual case, with h in the model being the wage rate) or, perhaps, amongst people with various time (hence saving) preferences (indexed by h), if the government is not committed to lifetime horizontal equity but to spot equality and control of the supply of savings, say [Seade (1977)]. But the important point here is that in any event (apart from the complications mentioned in footnote 29), the imposition of any set of distortions on an h -man's demands will actually yield positive redistributive gains only in so far as there are people on either side of this man, amongst which the redistribution is to take place.

One would like to know *how local* the endpoint result can be expected to be, in general or in particular cases. In the light of the above interpretation one may conjecture that the marginal tax at a given point of the income scale, for example, should be lower on efficiency grounds the more people there are *at* that point, for that means the distortion effect is augmented by the larger base on which it is being applied, and that it should be higher on distributional grounds the more people there are *beyond* that point towards either end of the tax schedule. Hence how well the endpoint result approximates the taxes required at other levels of income seems largely to depend on how rapidly the density $f(\cdot)$ falls down towards zero near its ends (or whether it does). Given the shape of income distributions usually encountered in practice (and if the underlying distribution of wages has a similar shape), this would suggest fairly low marginal taxes along a certain range of low incomes, which is in fact a feature of most current tax

schedules, having a certain level of income tax exemption. According to this, and contrary to the standard proposal, negative income taxes should basically amount (qualitatively speaking) to the replacement of current unemployment benefits by a fixed income subsidy, *not* to be coupled with an increase of the marginal taxes levied on the working poor (although revenue considerations of course require more taxes to be levied elsewhere, perhaps including a reduction in the exemption level). On the other hand, I would expect the zero marginal income tax at the top of the scale to be a relatively poor approximation to the way high incomes, on the whole, ought to be taxed. At each point of the long and thin upper tail of income (other than near the very top) few people will be facing the distortion imposed by a positive marginal tax, while thereby a substantial revenue (useful as such or for redistribution) will be levied from all those with incomes above that level.³⁰

Appendix (to section 2)

The purpose of this appendix is to derive the differential equation defining the family of Pareto-improved tax schedules of section 2.2. I assume: (i) that we already know or have estimates of the parameters of a unique utility function capturing consumers' preferences between consumption c and work done l , (ii) that consumption is noninferior (see footnote 8) and, for simplicity, (iii) that there is a constant basic wage rate (constant returns), and (iv) that θ is smooth and has no bunching at the bottom.

The marginal rate of substitution of c for y is given by

$$s \equiv s(c, y, w) = -\frac{u_l}{wu_c}, \quad (\text{A.1})$$

and by construction of $\hat{\theta}$, i.e., of $c = \hat{\theta}(y)$, at each point (c, y) on $\hat{\theta}$ we want

$$\hat{\theta}' \equiv \left. \frac{dc}{dy} \right|_{\hat{\theta}} = s(c, y, w). \quad (\text{A.2})$$

The problem is to find the relevant value of w to be used in (A.2). This must correspond to the wage rate of a man who, under θ , consumes a pair (c^0, y^0)

³⁰Mirrlees' (1971) asymptotic results seem to support (and can be explained by) this conjecture: he found rather high asymptotic marginal taxes (about 30–60%) for the thick and slowly falling Paretian upper tail, in contrast with fairly low marginal taxes at high incomes (15–20%, asymptotically zero) for the more rapidly falling lognormal (both for elasticities of substitution between consumption and leisure not greater than one).

which lies on the same 45° line as (c, y) . That is,

$$\begin{aligned} c - y &= c^0 - y^0 \\ &= \theta(y^0) - y^0 \\ &\equiv \xi(y^0). \end{aligned} \tag{A.3}$$

Assume $\theta' < 1$ at (c^0, y^0) , i.e. that $\xi'(y^0) < 0$ for all y^0 in the relevant domain of $\theta(\cdot)$. We can then invert this function:

$$\begin{aligned} y^0 &= \xi^{-1}(c - y) \\ &\equiv \varphi(c - y), \end{aligned} \tag{A.4}$$

which determines the pair $(c^0, y^0) \equiv (\theta(y^0), y^0)$ on θ that corresponds to each point (c, y) on $\hat{\theta}$. But on θ we must also have the tangency condition for each w (given smoothness of θ):

$$\theta'(y^0) = s(c^0, y^0, w), \tag{A.5}$$

which, by noninferiority of consumption, can be inverted for w :

$$w = w(y^0), \tag{A.6}$$

using $c^0 = \theta(y^0)$.

Substitution of (A.4) into (A.6) and of this in (A.2) finally yields the desired

$$\left. \frac{dc}{dy} \right|_0 = s(c, y, w[\varphi(c - y)]) \equiv g(c, y). \tag{A.7}$$

Let us now consider more explicitly the problem of uniqueness. If the function $g(c, y)$ is defined and continuous and has a bounded derivative with respect to c at (and in a neighbourhood of) a given point (\bar{c}, \bar{y}) , then an existence and uniqueness theorem for differential equations shows that one and only one integral curve for (A.7) passes through (\bar{c}, \bar{y}) [see, e.g., Sánchez (1968, p. 124)]. Assuming away uninteresting difficulties that might arise from the relation between w and y^0 on θ (we basically exclude the possibility that θ is such that no one chooses to earn income in an interior range of y^0 's), the conditions of the uniqueness theorem are met if and only if $\varphi' \equiv (\xi^{-1})'$ is bounded, which is to say if ξ' does not vanish, the positive-marginal-tax condition. Since θ is a particular solution for (A.7), it follows that any other solution which starts above θ at

the top, such as $\hat{\theta}$ in fig. 2, stays strictly above θ everywhere, so that the Pareto improvement is to the strict benefit of all w . But if at some point θ has $\theta' = 1$, g_c at that point is unbounded and the uniqueness theorem does not hold there. The highest such point is a singularity at which θ and every $\hat{\theta}$ merge together, remaining, by construction, the same from that point downwards.

The problem of selecting a starting point for $\hat{\theta}$, finally, amounts to finding the uppermost (and in fact best) allowable such point, a^* . This is easily done using the equation for the top man's 45° line (we have its slope, 1, and a point, a), which is tangent to an indifference curve of his at a^* .

References

- Atkinson, A.B., 1972, 'Maxi-min' and optimal income taxation, presented at the 1972 European Meeting of the Econometric Society in Budapest, mimeo.
- Atkinson, A.B., 1973, How progressive should income-tax be? in: M. Parkin, ed., *Essays in modern economics* (Longman, London) 90-109.
- Atkinson, A.B. and J.E. Stiglitz, 1974, Alternative approaches to the redistribution of income, for the International Seminar in Public Economics, Paris, January 1975, mimeo. Condensed version: The design of tax structure: Direct versus indirect taxation, *Journal of Public Economics* 6 (1976), 55-75.
- Feldstein, M.S., 1973, On the optimal progressivity of the income tax, *Journal of Public Economics* 2, 357-376.
- Flemming, J.S., 1976, *Optimal redistribution and social security funding*, Nuffield College, Oxford, mimeo.
- Gelfand, I.M. and S.V. Fomin, 1973, *Calculus of variations* (Prentice-Hall, Englewood Cliffs, NJ).
- Mirrlees, J.A., 1971, An exploration in the theory of optimum income taxation, *Review of Economic Studies* 38, 175-208.
- Mirrlees, J.A., 1976, Optimal tax theory: A synthesis, *Journal of Public Economics* 6, 327-358.
- Phelps, E.S., 1973, Taxation of wage income for economic justice, *Quarterly Journal of Economics* 87, 331-354.
- Pigou, A.C., 1946, *The Economics of welfare*, fourth edition (Macmillan, London) 87-97, 713-719.
- Sadka, E., 1976, On income distribution, incentive effects and optimal income taxation, *Review of Economic Studies* 43, 261-267.
- Sánchez, D.A., 1968, *Ordinary differential equations and stability theory: An introduction* (Freeman, San Francisco).
- Seade, J.K., 1977, *On optimal taxation theory*, D.Phil. thesis (to be submitted), Oxford.
- Stern, N.H., 1976, On the specification of models of optimum income taxation, *Journal of Public Economics* 6, no. 1/2, 123-162.