

ISSN 1471-0498



DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

**OPTIMUM INCOME TAXATION WHEN EARNINGS ARE
IMPERFECTLY CORRELATED WITH PRODUCTIVITY**

David L. Bevan

Number 101

May 2002

Manor Road Building, Oxford OX1 3UQ

Optimum income taxation when earnings are imperfectly correlated with productivity.

David L Bevan
Department of Economics
Manor Road
Oxford OX1 3UL
david.bevan@economics.ox.ac.uk

October 1997, Revised March 2002

Abstract

Probably the most enduring result in the theory of optimum income taxation is that, for a sufficiently thin upper tail to the skill distribution, the marginal tax rate should fall rather than rise with income. This paper shows that this result is highly sensitive to a very strong informational assumption, namely that earnings exactly reflect a worker's contribution to output. While the formal structure of the optimum problem is altered only slightly when earnings are allowed to be less than perfectly correlated with productivity, the shape of the optimum schedule is very sensitive to this relaxation. For high but imperfect correlation, optimum schedules look rather like those traditionally chosen by governments, with the marginal rate rising over high incomes and possibly U-shaped over the whole distribution.

Key words: Income taxation

JEL classification: H21, D82

Optimum income taxation when earnings are imperfectly correlated with productivity.

1. Introduction

This paper examines optimum income taxation when the informational problem is somewhat more extensive than that usually considered in this literature. It takes up a point made in Stiglitz, 1987, page 1021. “Views about the desirability of progressive taxation are often related to views concerning the extent to which differences in income are due to differences in effort, to differences in ability, or to differences in luck.” This observation might reflect the role of luck in weakening the view that those on high incomes have a moral right to them because these incomes are ‘earned’. However, this paper shows that the presence of even small amounts of luck radically alters optimum progressivity even in the standard welfarist set-up where considerations of rights do not arise.

Stiglitz (1982) examined the case where income was stochastic, but the government could still infer the individual’s ability from his income. He found that this could imply a marginal tax rate on high incomes of 100%, provided income was bounded above. He observed (page 235) “more generally, however, we will not be able to distinguish perfectly a low ability lucky individual from a high ability unlucky individual. This makes the design of the optimal tax structure ... far more difficult (and more interesting) ...”. The present paper sets out to address this issue.

Leaving aside a handful of very special cases, the optimum income tax problem is not analytically tractable, but requires numerical solution. This requires the structure of the problem to be specified very explicitly. Stern (1976, p.123) noted four ‘ingredients’ in this specification: an objective function, a preference relation or supply function for individuals, a skill structure and distribution, and a production relation. There is, however, in addition, the issue of incomplete information and how this is specified. This paper focuses on this fifth feature. It shows that quite mild variations in the informational assumptions, given the rest of the specification, can lead to a markedly different form for the optimum tax schedule. In particular, the common finding, both in numerical exercises (Mirrlees, 1971; Tuomala, 1990; Kanbur and Tuomala, 1994; Kanbur et al, 1994; Slemrod et al, 1994) and in asymptotic analyses (Mirrlees, 1971; Seade, 1977) that the marginal tax rate should eventually fall with income¹, is shown not to be robust to a relaxation in the usual informational set-up.

The structure of the paper is as follows. Section 2 outlines the informational approach adopted here, which essentially posits noise in the employer's assessment of a worker's contribution to output. In effect, the present model assumes more stringent informational problems than is usual in this literature. Section 3 briefly recapitulates the original Mirrlees model, and Section 4 then augments this to take the reduction in information into account. Section 5 introduces specific functional forms to permit numerical solution of the system. As regards Stern's ingredients, these are chosen to replicate those used in the original numerical illustration (Mirrlees, 1971 pages 200-

207). Given the nature of these assumptions, it is natural to model the measurement error as lognormally distributed and multiplicative. Section 6 concludes.

2. Informational assumptions

The usual informational assumption in this literature is related to the assumption on technology, that the individual's output is the deterministic product of his skill (ability) and his effort (hours worked), and that the output of individuals of different ability is homogeneous. The informational assumption is then that whereas output is perfectly observable, ability and effort are not separately observable at all. In this stylized setting, this is of no concern to the employer, who is only interested in the worker's output, and not in how it is produced; it does not matter, for example, whether a moderate output is the outcome of a high ability individual working a little, or a low ability individual working a lot. The government, by contrast, would like to be able to distinguish these cases, and tax by ability rather than by income, but this option is precluded by the separate non-observability of effort and ability.

The first point to notice about the construction of this principal-agent problem is how much it differs from the construction of the analogous principal-agent problem involving just the employer and worker. In the simplest form of the latter, ability differences between workers are suppressed. The problem arises because the output attributable to the individual can only be observed with error, either due to intrinsic difficulties in partitioning output, or because output is the joint product of the worker's efforts and an unobservable exogenous shock. The non-observability of effort may now become an issue for the employer, and the outcome is a wage that is imperfectly correlated with the worker's productivity². The question I wish to address here is what are the implications of this imperfect correlation between wages and productivity for the choice of the optimum tax schedule? To permit the analysis to focus on this core point, I assume that employers do not formulate complicated wage contracts, but simply pay workers *pro rata* for their imperfectly observed contribution to output³. For the same reason, the other features of the model closely follow those of Mirrlees' original paper.

3. Recapitulation of the Mirrlees model

While subsequent elaborations of the 1971 paper are cast in very general terms, the difficulty of obtaining clear cut and robust general results has meant that much weight has rested on the numerical illustrations provided in the original paper, and some limited variants of these provided by other authors, most notably Tuomala (1990). In this section, the bare bones of this special case are presented⁴; the following section adds the imperfect observability feature.

The fraction of the unit time endowment that an individual works is denoted by y . Individuals have identical preferences over consumption (x) of a composite good, and working time, $u = u(x, y)$, where u is a strictly concave, continuously differentiable function (strictly) increasing in x , (strictly) decreasing in y , defined for $x > 0$ and $1 > y \geq 0$. Individuals differ only in their skill or productivity per unit time. Skill is indexed by n with density $f(n)$, normalized so that:

$$\int_0^{\infty} f(n)dn = 1$$

Output is homogeneous. An 'n-person' produces an amount $z(n) = ny(n)$ and aggregate output is:

$$Z = \int_0^{\infty} ny(n)f(n)dn$$

The government is able to observe the individual's earnings exactly and these in turn exactly match his contribution to output, z , where the output price is set at unity. The government chooses the income tax schedule so that an individual with income z is constrained to consume no more than $c(z) = c(ny(n)) = x(n)$ after tax. In the absence of satiation, aggregate consumption is then:

$$X = \int_0^{\infty} x(n)f(n)dn$$

The government has a net revenue requirement \bar{R} so its budget constraint is $X + \bar{R} = Z$, where:

$$\int_0^n [ty(y) - x(t)] f(t)dt = R(n) \quad (1)$$

and $\bar{R} = R(\infty)$. It also has a welfare function that is increasing in individual utilities:

$$G = \int_0^{\infty} W(u^n) f(n)dn$$

where u^n is the utility of an n -person. This can be formulated as a maximum principle problem, with $y(n)$ as the control variable and $u(n)$ and $R(n)$ as the state variables. The first order condition for utility maximization is $u_x n c'(ny) + u_y = 0$. Hence, writing g for the derivative of u^n with respect to n :

$$g(u, y) = \frac{du^n}{dn} = u_n^n = u_x c' y = -\frac{yu_y}{n} \quad (2)$$

We now differentiate g given n invariant, yielding:

$$g_y = -\left[u_y + yu_{yy} - yu_{yx}u_y / u_x \right] / n$$

$$g_u = -yu_{yx} / nu_x$$

The Lagrangean for the government's problem is:

$$\int_0^{\infty} \{W(u^n) f(n) + \lambda(n) \left([ny(n) - x(n)] f(n) - \frac{dR(n)}{dn} \right) + \psi(n)(u_n^n - g)\} dn \quad (3)$$

Integration by parts yields the Hamiltonian:

$$H = \{W(u^n) + \lambda(n)[ny(n) - x(n)]\} f(n) - \psi(n)g(n) + \lambda_n R(n) - \psi_n u^n \quad (4)$$

with the first order conditions:

$$H_u = (W' - \lambda / u_x) f(n) - \psi_n - \psi g_u = 0 \quad (5)$$

$$H_y = \lambda(n + u_y / u_x) f(n) - \psi g_y = 0 \quad (6)$$

$$H_R = \lambda_n = 0 \quad (7)$$

Define:

$$v = \frac{1 + u_y / nu_x}{-(u_y + yu_{yy} - yu_{yx}u_y / u_x)} = \frac{\psi}{n^2 \lambda f(n)} \quad (8)$$

Then $\psi = n^2 \lambda f(n) v$,

$$\begin{aligned} \psi_n &= \left[\frac{2}{n} + \frac{\lambda_n}{\lambda} + \frac{f'}{f} + \frac{v_n}{v} \right] \psi \\ &= \left[\frac{2}{n} + \frac{f'}{f} + \frac{v_n}{v} \right] \psi \end{aligned} \quad (9)$$

and:

$$\left[\frac{2}{n} + \frac{f'}{f} + \frac{v_n}{v} - \frac{yu_{yx}}{nu_x} \right] v = \frac{(W' - \lambda / u_x)}{n^2 \lambda} \quad (10)$$

Equations (2) and (10) can be written as:

$$u_n = -yu_y / n \quad (11)$$

$$v_n = -\frac{v}{n} \left[2 + nf' / f - yu_{yx} / u_x \right] + (W' / \lambda - 1 / u_x) / n^2 \quad (12)$$

This pair of first order ordinary differential equations can then be solved for the optimum schedule, once specific functional forms for u , W , and f are chosen.

Mirrlees also showed⁵ that there exists a number $n_0 \geq 0$ such that:

$$y^n = 0(n \leq n_0), \quad y^n > 0(n > n_0)$$

and argued that $n_0 > 0$ for distributions of the type “usually considered by economists” (Mirrlees, 1971 p.185). In other words, it would usually be (second-best) optimal to have some of the population idle. n_0 provides a convenient method for establishing initial values for the solution.

4. Incorporating imperfectly observed productivity

4.1 Representation

The simplest way of introducing this feature is to postulate a bivariate distribution of true and apparent ability. It is natural to assume that the two marginal distributions take the same form. In the original numerical illustration true (the only) ability was assumed to be distributed in the population lognormally, $n \sim \Lambda(\mu_n, \sigma_n)$ where $\ln \mu_n$ and σ_n are the mean and standard deviation of the associated normal distribution, i.e. n has median μ_n and mean $\mu_n e^{\sigma_n^2/2}$, and we maintain this assumption for the marginal true distribution. Now consider a worker i , who has true ability n_i and works y_i . His true contribution to output is therefore $q_i = n_i y_i$, but this is perceived by the employer as $z_i = m_i y_i$ where $z_i = q_i \varphi$ and φ is an error term. It follows that the same relation holds in the unobservables, so that $m_i = n_i \varphi$. We assume that the error $\varphi \sim \Lambda(\mu_\varphi, \sigma_\varphi)$, and that it may be correlated with n , with $\rho_{n\varphi}$ as the correlation coefficient in the associated bivariate normal distribution. Then it follows from the properties of the normal distribution that:

$$\sigma_m^2 = \sigma_n^2 + \sigma_\varphi^2 + 2\rho_{n\varphi}\sigma_n\sigma_\varphi \quad (13)$$

$$\rho_{mn} = (\sigma_n + \rho_{n\varphi}\sigma_\varphi) / \sigma_m = \frac{\sigma_n + \rho_{n\varphi}\sigma_\varphi}{\sqrt{\sigma_n^2 + \sigma_\varphi^2 + 2\rho_{n\varphi}\sigma_n\sigma_\varphi}} \quad (14)$$

It is worth considering some special cases of this relation. Evidently, if $\rho_{n\varphi} = 1$, then so does ρ_{mn} , and the error φ is reduced to a scaling factor. The ranking of individual workers by productivity is now exactly right, but the dispersion is exaggerated in the ratio $(1 + \sigma_\varphi / \sigma_n)^2$. In all other cases, $|\rho_{n\varphi}| < 1 \Rightarrow |\rho_{mn}| < 1$. What is more, $\rho_{mn} > 0$ even when $\rho_{n\varphi} < 0$ provided $-\rho_{n\varphi} < \sigma_n / \sigma_\varphi$ and always when $\rho_{n\varphi} > 0$. Provided the variance of $\ln \varphi$ is less than that of $\ln n$, so that the measurement error is no larger than the true variation in the object measured, as seems plausible, this condition will hold. Hence a positive but imperfect correlation between true and apparent ability would seem to be robust to any plausible parameterization of the error.

We have the marginal distribution with $m \sim \Lambda(\mu_m, \sigma_m)$. Since we shall suppose that the parameters of the true skill distribution (μ_n, σ_n) and those of the error $(\sigma_\varphi, \rho_{n\varphi})$ are given, this only leaves μ_m to be determined. We make the natural assumption that the employer is accurately informed about total (average) product. We also assume that each employer has a balanced labour force, and is not specialized in some part of the ability distribution. Then if working time were uniform, $y(m) = y$, all m , this would imply the restriction that the means of the two distributions coincide,

$$\mu_m = \mu_n \mu_\varphi = \mu_n e^{-(\sigma_\varphi^2 + 2\rho_{n\varphi}\sigma_n\sigma_\varphi)/2}.$$

However, we are interested in situations where labour supply will not in general be uniform across worker types, and where it will respond differentially to changes in the tax regime. The relation between skill and earnings then becomes more delicate, and it is useful to separate the apparent skill index m of a worker from the wage he is paid, $z(m) = wmy(m)$, where w is common across all workers. We can then treat m as a characteristic of the worker which describes his apparent skill (and wage per unit time) relative to other workers, but does not necessarily equal his real product wage. We can choose μ_m arbitrarily, preserving the values m / μ_m . In what follows, we choose $\mu_m = e^{(\sigma_n^2 - \sigma_m^2)/2}$ since that equates average per period values of true and apparent skill.

4.2 Incorporation (1)

In place of an n -worker, we now have an (m, n) worker, where m is his apparent and n his true ability. Consider an (m, \cdot) worker: he is treated as a Mirrlees m -worker, and he makes the labour supply decision ($y(m)$) and obtains the consumption ($x(m)$) of a Mirrlees m -worker. However, there is now a distribution of product contributed by the group of workers classified as (m, \cdot) workers which is $y(m)$ times the conditional distribution of n given m . Given that $m, n \sim \Lambda(\mu_n, \mu_m, \sigma_n, \sigma_m, \rho_{mn})$, and writing $h(m)$ for the expected value of this conditional distribution, we have $n | m \sim \Lambda(\mu_{n|m}, \sigma_{n|m})$

where $\mu_{n|m} = \mu_n (m / \mu_m)^{\rho_{mn}\sigma_n/\sigma_m}$ and $\sigma_{n|m}^2 = \sigma_n^2 (1 - \rho_{mn}^2)$. Hence

$h(m) = \mu_n (m / \mu_m)^{\rho_{mn}\sigma_n/\sigma_m} e^{\sigma_n^2(1-\rho_{mn}^2)/2}$. Incorporating the condition $\mu_m = e^{(\sigma_n^2 - \sigma_m^2)/2}$, we obtain:

$$\begin{aligned} h(m) &= \mu_m e^{(\sigma_m^2 - \sigma_n^2)/2} (m / \mu_m)^{\rho_{mn}\sigma_n/\sigma_m} e^{\sigma_n^2(1-\rho_{mn}^2)/2} \\ &= \mu_m^{(1-\rho)} e^{(1-\rho^2)\sigma_m^2/2} m^\rho \end{aligned} \tag{15}$$

where we define $\rho = \rho_{mn}\sigma_n / \sigma_m$. Substituting from equations (13) and (14) we obtain:

$$\rho = \frac{\sigma_n^2 + \rho_{n\varphi}\sigma_n\sigma_\varphi}{\sigma_n^2 + \sigma_\varphi^2 + 2\rho_{n\varphi}\sigma_n\sigma_\varphi}$$

Provided there is any error at all ($\sigma_\varphi^2 \neq 0$), then $\rho < 1$ unless $\rho_{n\varphi}$ is negative and $|\rho_{n\varphi}| \geq \sigma_\varphi / \sigma_n$. While alternatives have been discussed for completeness, the most natural assumption would be that the measurement error is uncorrelated with ability. In that case, we get the simple relation $\rho = \sigma_n^2 / (\sigma_n^2 + \sigma_\varphi^2) = \sigma_n^2 / \sigma_m^2$. The associated condition, $E[m | n] = n$, all n , is seen to imply that $E[n | m] \equiv h(m) \neq m$ in general. However, even for biased errors, within any plausible range of bias, the implication of this discussion is that $h(m) \propto m^\rho$ with $\rho < 1$ so that $h(m)/m$ is decreasing in m .

4.3 Incorporation (2)

Now consider how to recast the original model to take this effect into account. We could in principle continue to integrate over the variable n , but this would be cumbersome, since workers who share the same value of true ability are now treated differently by their employers and hence by the tax authorities. In consequence, they face different budget constraints and will generally behave differently. This difficulty does not arise if we integrate instead over the variable m . The behaviour of an (m, n) worker, as already noted, is indistinguishable from that of a Mirrlees m -worker. Workers who share the same apparent ability are treated the same, face the same budget constraint, and behave the same. The characterization of social welfare is similarly unaffected if we replace n with m . The only change is that aggregate output attributable to workers with apparent skill $t \leq m$ is:

$$Q(m) = \int_0^m h(t)y(t)f(t)dt$$

where $h(t)$ is given by equation (15). Making the associated changes to the government's maximization problem, we need to replace the government's budget constraint (1) by:

$$\int_0^m [h(t)y(t) - x(t)]f(t)dt = R(m)$$

The only changes to the Hamiltonian (4) are the switch for the variable of integration from n to m , and the replacement of the term in square brackets by $[h(m)y(m) - x(m)]$. The first order conditions are undisturbed, with the exception of (6) which becomes:

$$H_y = \lambda(h(m) + u_y / u_x) f(m) - \psi g_y = 0$$

This alters the definition of v :

$$v = \frac{h(m)/m + u_y / mu_x}{-(u_y + yu_{yy} - yu_{yx}u_y / u_x)}$$

However, once this alteration is accounted for, the form of the two differential equations (11) and (12) remains the same.

The other way in which the changed specification alters the solution is via the definition of gross income (z). Since output and income are both defined in real terms, we must have $Q = Z$. However, for a predetermined value of μ_m ,

$$\int_0^{\infty} my(m)f(m)dm \neq \int_0^{\infty} h(m)y(m)f(m)dm$$

unless $y(m)=y$, or $\rho = 1$. We proceed by defining a wage factor w :

$$w = \frac{\int_0^{\infty} h(m)y(m)f(m)dm}{\int_0^{\infty} my(m)f(m)dm}$$

We now have $Z = w \int_0^{\infty} my(m)f(m)dm = Q$ as required⁶, and that the individual (m, \cdot)

worker has gross income $z(m)=wmy(m)$. The upshot is that workers are paid in proportion to their perceived productivities, and that the economy-wide average wage equals the economy-wide average product. The government is effectively choosing the consumption schedule directly in light of real output and revenue requirements. We thus have that the worker pays tax of $wmy(m)-x(m)$; writing r for the share of government revenue in total purchasing power, we have that:

$$r = R / Z = 1 - X / Z$$

$$= 1 - \frac{\int_0^{\infty} x(m)f(m)dm}{w \int_0^{\infty} my(m)f(m)dm}$$

For the purpose of comparative tax analysis, we here choose to maintain r constant.

4.4 Interpretation

It may seem that the change in the optimum problem when imperfect correlation between earnings and productivity is introduced is rather slight. This is certainly true of its structure. However, the question of interest here is the numerical one, as to how sensitive the optimum tax schedule is to plausible values of this correlation. Before turning to computation, it is worth considering the qualitative nature of the change. The central problem in optimum taxation is the trade-off between equity and efficiency. Highly able workers receive high incomes, so are desirable targets for high, redistributive taxes. The problem is that they are paid highly precisely because they are highly productive. Hence disincentive effects, leading to a reduction in their labour supply may be socially very costly. At the other end of the spectrum, redistribution toward the low paid is desirable from the equity viewpoint, and the associated incentive effects may be relatively unimportant, because these workers are

very unproductive anyway. These features help to explain two characteristics of the usual optimum calculations. First, marginal tax rates are low, and typically falling at the top of the distribution, and second, there may be a group at the bottom of the distribution for whom it is socially optimal that they do no work.

If we now introduce imperfect correlation between earnings and productivity, the nature of this trade-off is shifted, and is to some extent less stark. Highly paid workers are paid, on average, more than they are worth; the low paid are paid less than they are worth. The potential equity gains from high taxes on highly paid workers are unchanged, but the production losses from a unit reduction in their labour supply are reduced. Indeed, if the correlation were very low, the highly paid would be receiving substantial economic rents which are usually seen as appropriate targets for high taxes. However, the incentive effects on the low paid are now more serious; while the poor are in one sense more deserving in this case, they are harder to compensate. It should be clear from this discussion that imperfect correlation of this type is likely to have complex effects, and these may not vary in any simple way with the size of the correlation coefficient. In particular, it is not clear whether a fall in the perceptibility of productivity would make the optimal tax system more or less redistributive.

5. Numerical illustrations

5.1 Assumptions

These assumptions follow closely those chosen for the numerical illustrations in the 1971 paper. As already noted, the distribution of n was there assumed to be lognormal. Mirrlees took $\mu_n = e^{-1}$ and, for five of the six calculations reported, $\sigma_n = 0.39$; we restrict attention to these values, but now for the distribution of m . Utility is assumed to take the Cobb-Douglas form, $u = x(1 - y)^{1/\alpha}$ and the welfare function the form:

$$G = \int_0^{\infty} [(u^{1-\gamma} - 1)/(1 - \gamma)] f(n) dn$$

In the calculations, α and γ are both set at 1⁷. The net revenue requirement was set at rather arbitrary amounts in the original calculations, presumably as a consequence of the solution procedure. This was to choose trial values of λ and n_0 , and iterate on n_0 until a sufficiently good fit was obtained. The ratio of aggregate consumption to aggregate production, X/Z and hence the net revenue requirement at this solution depend on the trial value of λ originally chosen. In the present calculations, a second round of iterations are undertaken, on λ , to maintain the revenue requirement constant between experiments.

5.2 Results

Since the interest of the present exercise lies in the shape of the tax function, the results are presented graphically in Figure 1. In panel a, the tax is purely redistributive, while in panel b it must also meet a net revenue requirement equal to

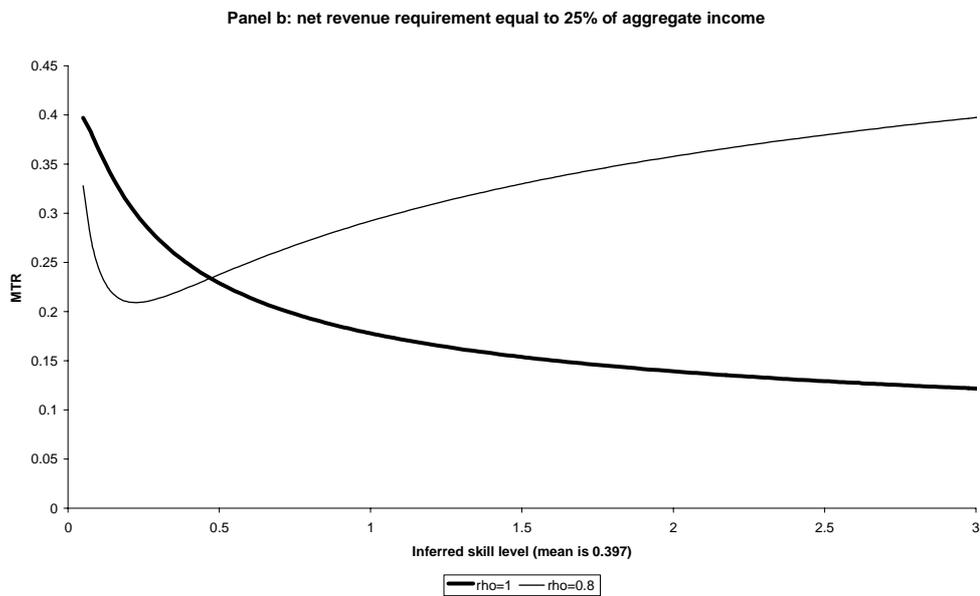
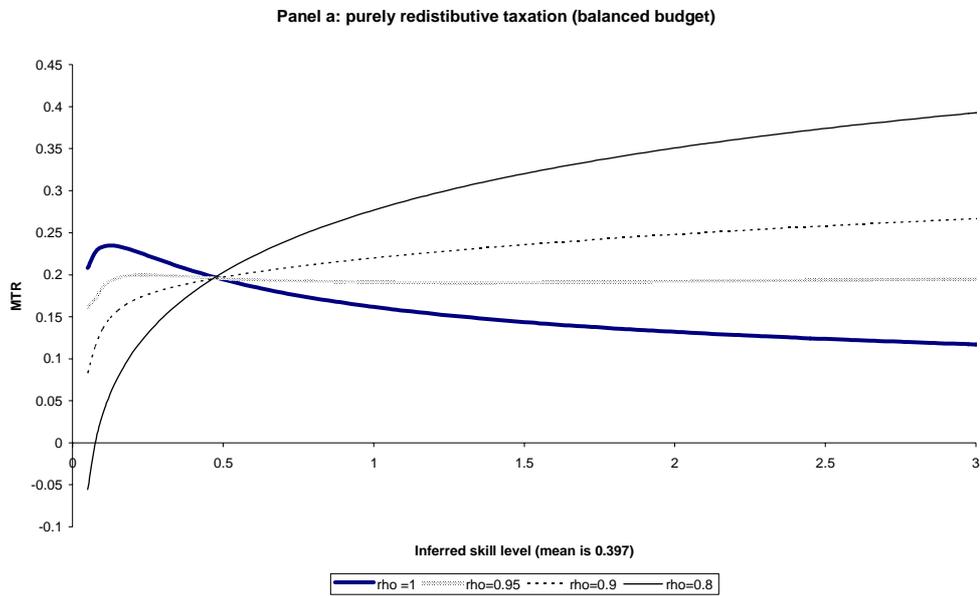
25% of income. Panel a graphs the marginal tax rate for four degrees of correlation between true and inferred skill⁸. The bold line plots the tax rate when the correlation is perfect, i.e. this is the case considered by Mirrlees. The rate rises slightly for very low skill levels (above the level below which individuals do no work), but then falls rather slowly. It was this type of calculation that led Mirrlees to conclude (i) that asymptotic results were rather misleading and (ii) that a linear tax might not be a bad approximation to the optimal tax. The remaining three plots in the panel explore what happens as we progressively relax the link between what workers are paid and what they produce. Even when the correlation is as high as 0.95⁹, the tax schedule undergoes a substantial shift, becoming slightly u-shaped, but very close to being stationary. As the correlation falls to 0.9, the tax schedule now shows a rather rapid rise at low skill levels, and continues to rise steadily across the upper range. With a correlation of 0.8, this rise becomes quite marked, and not dissimilar from the pattern of conventional income tax schedules adopted in practice. Also, at low inferred skill levels, it is now optimal for the marginal tax rate to be negative. The disincentive effects of transferring income to the low paid are now becoming serious, since these workers are on average substantially more productive than they appear to be from their wages¹⁰.

Panel b plots the schedules for correlations of 1.0 and 0.8 when there is a positive revenue requirement. The interesting feature here is that the displacement of the marginal schedules is rather slight for high skill levels, but shows a very marked increase at low levels. With a correlation of 0.8, the imposition of the net revenue requirement makes the optimum schedule u-shaped.

6. Conclusion

Once we allow that wages are not perfectly aligned with productivity, the shape of optimum income tax schedules may be altered drastically. Even if the wage-productivity fit is good, it is likely that the marginal tax rate should rise with income, at least in the upper ranges, even in the lognormal case where the conventional treatment indicates a declining rate. For more realistic specifications of the skill distribution, with a thicker upper tail, this conclusion is reinforced.

Figure 1: Marginal tax rate schedules (MTR) for different degrees of correlation between observed and true productivity (ρ)



References

- Bevan, D.L. "Principals and Agents when Ability Differs and is Imperfectly Observed." mimeo, Oxford, October 1997
- Diamond, P.A. "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *American Economic Review*, 88 (1998): 83-95
- Ebert, U. "A Re-examination of the Optimal Nonlinear Income Tax." *Journal of Public Economics* 49 (1992): 47-73
- Kanbur, R., and M. Tuomala. "Inherent Inequality and the Optimal Graduation of Marginal Tax Rates." *Scandinavian Journal of Economics* 96 (1994): 275-282
- Kanbur, R., M. Keen and M. Tuomala. "Optimal Non-Linear Income Taxation for the Alleviation of Income-Poverty." 1994, 38, pp. 1613-1632
- Mirrlees, J.A. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (1971): 175-208
- Mirrlees, J.A. "Optimal Tax Theory: A Synthesis." *Journal of Public Economics* 6 (1976): 327-358
- Mirrlees, J.A. "The Theory of Optimal Taxation." in K.J. Arrow and M.D. Intriligator, eds, *Handbook of Mathematical Economics Vol. III*. North-Holland, Amsterdam 1986: chapter 24, 1197-1249
- Seade, J. "On the Shape of Optimal Tax Schedules." *Journal of Public Economics* 7 (1977): 203-235
- Seade, J. "On the Sign of the Optimum Marginal Income Tax." *Review of Economic Studies* 49 (1982): 637-643
- Slemrod, J., S. Yitzhaki, J. Mayshar, and M. Lundholm. "The Optimal Two-Bracket Linear Income Tax." *Journal of Public Economics* 53 (1994) 269-290
- Stern, N.H. "On the Specification of Models of Optimum Income Taxation." *Journal of Public Economics* 6 (1976): 123-162
- Stiglitz, J.E. "Self-Selection and Pareto Efficient Taxation." *Journal of Public Economics*, 17 (1982): 213-240
- Stiglitz, J.E. "Pareto Efficient and Optimal Taxation and the New New Welfare Economics." In Auerbach, A. J. and M. Feldstein, eds. *Handbook of Public Economics Volume 2*. North Holland, Amsterdam, 1987.
- Tuomala, M. *Optimal Income Tax and Redistribution*. Clarendon Press, Oxford, 1990.

Endnotes

¹ Mirrlees (1971) demonstrated that this would not necessarily hold for distributions with a Pareto upper tail, but this point has tended to be overlooked. See however Diamond 1998.

² This is true if the employer writes a second best optimal contract, trading off the worker's risk aversion against the disincentive effect of insuring him. But it is also true of any other feasible contract, including the simplest one that pays the worker for the observed output, leaving him to bear all the risk.

³ In a companion paper, Bevan (1997), I examine the case when employers and government are tackling their principal agent problems simultaneously, restricting attention to loglinear wage contracts and loglinear income tax schedules. In the present context, the most natural assumption is that employers cannot attribute output to individual workers, but can observe effort accurately, and have a shrewd but not perfect idea of workers' abilities. There is then no viable contract other than one that is proportional to this perceived contribution.

⁴ We assume differentiability of the relevant functions, and restrict attention to the first order conditions. For all the schedules computed in this paper, it turns out that the second order condition is strictly satisfied. See Ebert, 1992.

⁵ Mirrlees, 1971, proposition 1, p.177.

⁶ Note that the scale of the corrections is likely to be very small for reasonable values of the correlation, i.e. w is typically very close to 1.

⁷ These components were organised slightly differently in the original paper, with $u = \alpha \ln x + \ln(1 - y)$ and $W = -e^{-\beta u} / \beta$. For the calculations, α was set at 1 and β took the alternative values 0 and 1. These correspond to the forms in the text with α again set at 1 and γ set at 1 and 2 respectively.

⁸ The tax rates are plotted against the inferred skill level rather than income since it is the former that is invariant between cases. However the general shape of the schedules is very similar if plotted against income.

⁹ Since the "size" of these correlations is not intuitive, it is helpful to consider the implied relation between relative wages and productivity across some range of the distribution. For example, the group of workers paid twice the average wage will (on average, across the two groups) be "overpaid" relative to those paid the average wage by 3% if $\rho = 0.95$; by 6% if $\rho = 0.9$; and by 12% if $\rho = 0.8$.

¹⁰ Indeed for lower correlations than those displayed, the marginal tax schedule becomes u-shaped, with the bottom of the u being substantially negative.