

Means-testing and tax rates on earnings*

Mike Brewer, IFS
Emmanuel Saez, UC Berkeley
Andrew Shephard, UCL and IFS

April 8, 2008

PRESS COPY
Embargoed until 0001am,
Monday 21st April 2008

*This paper has been prepared for the Mirrlees Review - Reforming the Tax System for the 21st Century, <http://www.ifs.org.uk/mirrleesreview/>. Mike Brewer, Institute for Fiscal Studies, mike_b@ifs.org.uk, Emmanuel Saez, saez@econ.berkeley.edu, University of California, Department of Economics, 549 Evans Hall #3880, Berkeley, CA 94720, Andrew Shephard, Institute for Fiscal Studies, andrew.s@ifs.org.uk. We thank Stuart Adam, Tony Atkinson, Kate Bell, Richard Blundell, Hilary Hoynes, Paul Johnson, Guy Laroque, Costas Meghir, James Mirrlees, Robert Moffitt, James Poterba and numerous conference participants for helpful comments and discussions. Saez acknowledges financial support from the National Science Foundation grant SES-0134946. The Survey of Personal Incomes, the Labour Force Survey, and Family Expenditure Survey, the Family Resources Survey and the General Household Survey datasets are crown copyright material, and are reproduced with the permission of the Controller of HMSO and the Queen's Printer for Scotland. The SPI, LFS and GHS data-sets were obtained from the UK Data Archive, and FRS from the Department for Work and Pensions, and the FES from the Office for National Statistics. None of these government department nor the UK Data Archive bears any responsibility for their further analysis or interpretation.

The setting of income tax rates and the generosity and structure of income support programmes generate substantial controversy among policy-makers and economists. At the centre is a trade-off between the goals of equity and efficiency: governments want to transfer resources from the rich to the poor; on the other hand, such transfers reduce people's incentive to work.

The key insight from the standard "optimal income tax model" developed by James Mirrlees is that marginal rates of tax and benefit withdrawal should be higher when people's choices of how much to work are relatively unresponsive to them and when the government is relatively keen to redistribute resources from rich to poor. Furthermore, the government should apply high marginal rates at points in the earnings distribution where there are few taxpayers relative to the number of taxpayers who have earnings exceeding this amount. Using data on the UK earnings distribution, we show that the optimal structure of marginal rates in this simplified model has a U-shaped pattern, with high marginal rates imposed on high and low earners and lower marginal rates on those in the middle. We show how this structure changes as both the assumed responsiveness of hours of work and the government's assumed preferences for redistribution vary.

The way that incomes have responded to the large changes in top marginal tax rates over the past 40 years suggests that if the richest 1% see a 1% fall in the proportion of each additional pound of earnings that is left after tax, then the income they report will rise by less than half that - only 0.46%. Although a tentative estimate, this suggests that the government would maximise the revenue it collects by imposing an overall marginal rate on the highest earners of 56.6%, very close to the 53.0% currently charged in the UK (including income tax, national insurance contributions and indirect taxes). So there does not seem a powerful case for increasing the income tax rate on the very highest earners, even on redistributive grounds - it would not generate much if any extra revenue to transfer to the less well off.

When the optimal tax model is enriched by allowing individuals to respond to taxes and transfers by deciding whether or not to work, as well as how hard, then the optimal structure of marginal rates changes dramatically. In particular, when the decision whether to work becomes relatively more important to the decision about how much to work, then marginal rates and the proportion of gross income taken in tax and withdrawn benefits when people enter work should be set low (and perhaps even negative) for potential low earners rather than set high as the standard model suggests. We also discuss how the design of taxes and benefits affecting an individual should be affected by the presence of a co-resident partner or dependent children, although it is difficult to reach definitive conclusions. We argue that the practical operation of benefits and tax credits for low-income families is important and that they would be of greatest help to beneficiaries if they were assessed over short periods and paid promptly without retrospective adjustment.

These insights from optimal tax theory are contrasted with the work incentives inherent in the current UK tax and transfer system. Four key deficiencies are identified:

1. The amount of gross income taken in tax and withdrawn benefits when people enter work at low earnings is too high: for most groups it is close to 100% before individuals are entitled to the working tax credit, and they remain high even with it.

2. The marginal rate of 73.4% that many low to moderate earners face when having tax credits withdrawn is likely to be above the optimum rate even if people's decision to work a little harder is relatively unresponsive.
3. Housing Benefit, the main means-tested programme through which the government helps people on relatively low incomes with their housing costs has an extremely high withdrawal rate. This exacerbates the problem of undesirably high marginal rates. It is also hard to administer and is not claimed by many working families entitled to it.
4. While the system for administering income tax and payroll tax in the UK is simple and efficient, tax credits, housing benefit and council tax benefit are all burdensome to claim, relatively expensive for the government to administer, and prone to significant fraud and error.

Given this diagnosis, we suggest a set of changes to the existing tax and transfer structure that could be made immediately based on the lessons from our analysis. Our package of "immediate reforms" involves:

- Increasing the amount people can earn before they have means-tested benefits withdrawn. This would increase the financial gain to entering work at low earnings;
- Increasing the amount that second earners can earn before a family's tax credits are withdrawn. This would improve the financial incentive for a second earner to enter work, especially if they have children;
- Reducing the rate at which child and working tax credits are withdrawn with every extra pound earned;
- Targeting increases in working tax credit on groups other than lone parents.

This would cost around £9 billion per year. If it had to be financed from within the income tax and benefit system, the money could be raised by cutting child benefit and/or increasing the basic rate of income tax. Neither would undo the objectives of the reform package to improve work incentives, although both would pose big political challenges.

We also suggest a more radical and comprehensive plan for reforming the UK household tax and transfer system that attempts to deal not only with these work incentive issues, but also the administrative failings that we identify. Our plan replaces the existing piece-meal transfer programmes for low-income families (income support, working and child tax credits, housing benefit and council tax benefit) with a single Integrated Family Support (IFS) programme which provides stronger and simpler incentives for work at the bottom, reduces compliance costs for families, and is means-tested by employers' withholding from earnings in the same way as National Insurance contributions. We show how, after including an assessment of the behavioural responses, the IFS manages to redistribute more income with minimal impact on total earnings and total net tax revenue, by targeting net tax cuts where incentives to work are currently at their weakest.

1 Summary and key findings

The setting of income tax rates and the generosity and structure of income support programmes generate substantial controversy among policy-makers and economists. At the centre is a trade-off between the goals of equity and efficiency: governments want to transfer resources from the rich to the poor; on the other hand, such transfers reduce people's incentive to work.

The key insight from the standard "optimal income tax model" developed by James Mirrlees is that marginal rates of tax and benefit withdrawal should be higher when people's choices of how much to work are relatively unresponsive to them and when the government is relatively keen to redistribute resources from rich to poor. Furthermore, the government should apply high marginal rates at points in the earnings distribution where there are few taxpayers relative to the number of taxpayers who have earnings exceeding this amount. Using data on the UK earnings distribution, we show that the optimal structure of marginal rates in this simplified model has a U-shaped pattern, with high marginal rates imposed on high and low earners and lower marginal rates on those in the middle. We show how this structure changes as both the assumed responsiveness of hours of work and the government's assumed preferences for redistribution vary.

The way that incomes have responded to the large changes in top marginal tax rates over the past 40 years suggests that if the richest 1% see a 1% fall in the proportion of each additional pound of earnings that is left after tax, then the income they report will rise by less than half that - only 0.46%. Although a tentative estimate, this suggests that the government would maximise the revenue it collects by imposing an overall marginal rate on the highest earners of 56.6%, very close to the 53.0% currently charged in the UK (including income tax, national insurance contributions and indirect taxes). So there does not seem a powerful case for increasing the income tax rate on the very highest earners, even on redistributive grounds - it would not generate much if any extra revenue to transfer to the less well off.

When the optimal tax model is enriched by allowing individuals to respond to taxes and benefits by deciding whether or not to work, as well as how hard, then the optimal structure of marginal rates changes dramatically. In particular, when the decision whether to work becomes relatively more important to the decision about how much to work, then marginal rates and the proportion of gross income taken in tax and withdrawn benefits when people

enter work should be set low (and perhaps even negative) for potential low earners rather than set high as the standard model suggests. We also discuss how the design of taxes and benefits affecting an individual should be affected by the presence of a co-resident partner or dependent children, although it is difficult to reach definitive conclusions. We argue that the practical operation of benefits and tax credits for low-income families is important and that they would be of greatest help to beneficiaries if they were assessed over short periods and paid promptly without retrospective adjustment.

These insights from optimal tax theory are contrasted with the work incentives inherent in the current UK tax and benefit system. Four key deficiencies are identified:

1. The amount of gross income taken in tax and withdrawn benefits when people enter work at low earnings is too high: for most groups it is close to 100% before individuals are entitled to the working tax credit, and they remain high even with it.
2. The marginal rate of 73.4% that many low to moderate earners face when having tax credits withdrawn is likely to be above the optimum rate even if people's decision to work a little harder is relatively unresponsive.
3. Housing Benefit, the main means-tested programme through which the government helps people on relatively low incomes with their housing costs has an extremely high withdrawal rate. This exacerbates the problem of undesirably high marginal rates. It is also hard to administer and is not claimed by many working families entitled to it.
4. While the system for administering income tax and national insurance contributions in the UK is simple and efficient, tax credits, housing benefit and council tax benefit are all burdensome to claim, relatively expensive for the government to administer, and prone to significant fraud and error.

Given this diagnosis, we suggest a set of changes to the existing tax and benefit structure that could be made immediately based on the lessons from our analysis. Our package of "immediate reforms" involves:

- Increasing the amount people can earn before they have means-tested benefits withdrawn. This would increase the financial gain to entering work at low earnings;

- Increasing the amount that second earners can earn before a family's tax credits are withdrawn. This would improve the financial incentive for a second earner to enter work, especially if they have children;
- Reducing the rate at which child and working tax credits are withdrawn with every extra pound earned;
- Targeting increases in working tax credit on groups other than lone parents.

This would cost around £9 billion per year. If it had to be financed from within the income tax and benefit system, the money could be raised by cutting child benefit and/or increasing the basic rate of income tax. Neither would undo the objectives of the reform package to improve work incentives, although both would pose big political challenges.

We also suggest a more radical and comprehensive plan for reforming the UK household tax and benefit system that attempts to deal not only with these work incentive issues, but also the administrative failings that we identify. Our plan replaces the existing piece-meal benefits for low-income families (income support, working and child tax credits, housing benefit and council tax benefit) with a single Integrated Family Support (IFS) programme which provides stronger and simpler incentives for work at the bottom, reduces compliance costs for families, and is means-tested by employers' withholding from earnings in the same way as National Insurance contributions. We show how, after including an assessment of the behavioural responses, the IFS manages to redistribute more income with minimal impact on total earnings and total net tax revenue, by targeting net tax cuts where incentives to work are currently at their weakest.

2 Introduction

The setting of income tax rates, and the generosity and structure of income support (or transfer) programmes generate substantial controversy among policy-makers and economists. At the centre is an equity-efficiency trade-off. On the one hand, governments value redistribution, and so want to transfer resources from the rich to the poor, usually by taxing the incomes of the rich and subsidising the incomes of the poor. On the other hand, this redistribution is generally costly in terms of economic efficiency because of the disincentive effects of taxes and transfers (we explain this in more detail in section 3). The costs arise for two reasons : first,

raising income taxes may weaken the labour supply and entrepreneurship incentives of middle and high income individuals who face the taxes. Second, income transfer programmes may weaken the labour supply incentives of their recipients. These two responses can substantially raise the cost of improving the living standards of low income families.

The goal of this chapter is to provide an overview of the way economists think about the design of taxes and benefits affecting households, and to apply the lessons from this literature to the design of the UK tax and benefit system.

In economics research, the problem of designing taxes and benefits is tackled in two steps. The first step is a positive analysis, where economists develop models of individual behaviour to understand how individuals' work decisions respond to taxes and benefits. The central part of the positive analysis is the empirical estimation of models of individual behaviour, and there is a very broad literature that tries to estimate the size of the behavioral responses to taxes and benefits.¹

The second step is the normative analysis, or optimal policy analysis. Using models developed in the positive analysis, the normative analysis investigates what structure and size of the tax and benefit system would best meet a given set of policy goals; following Mirrlees (1971), economists call this line of research "optimal tax theory". Despite its name, optimal tax theory concerns itself just as much with the design of benefits as it does the setting of income tax rates: one of the key concepts of optimal tax theory is that of a net tax function, whereby people with high incomes pay some of that income in positive taxes to the government, and people with a low income receive money from the government (by paying negative taxes); no conceptual distinction is made between net recipients from and net contributors to the state's finances.²

At its heart, optimal tax theory says that the two desirable features of a tax and benefit system are that it be fair, and that it minimise disincentive effects.³ But the problem of having two desirable features is that one has to know how much weight to give to each. For example, a poll tax (under which all individuals have to pay the same level of tax) might have

¹The way that these models are estimated, and the key insights, are summarised in Meghir and Phillips, this volume.

²One difference between the tax system and the transfer system is that the former is usually cheaper to administer, and these distinctions can be reflected in more complicated optimal tax models.

³More complicated models can allow for other desirable features: one might be that a tax and transfer system is cheap to administer; Shaw et al, this volume, consider how this alters optimal tax models.

no disincentive effects, but is rather unfair to those on low incomes. As Heady (1993, p17) says, “the approach of the optimal tax literature is to use economic analysis to combine these criteria into one”. It does this by saying that the objective of the government when designing the tax and benefit system should be to maximise social welfare (subject to a need to raise a certain amount of revenue). Precisely how social welfare is expressed is not relevant at this stage, but the idea is that it reflects in a single index (or number) the desire both to have the economy as large as possible (because this directly increases people’s well-being) but also to have the income distributed as equally as possible. The expression for social welfare precisely quantifies the trade-off between these two desiderata: returning to the previous example of an economy with only a poll-tax, replacing that with an income tax which raised the same amount of money would give a more equal distribution of income, but - if there any disincentive effects to taxation - a smaller economy.

The normative analysis is crucial for policy-making because it shows how taxes and benefits should be designed in order to best attain the goals of the policy-maker. In particular, the normative analysis allows one to assess separately how changes in the redistributive criterion of the government, and changes in the size of the behavioural responses to taxes and transfers, affect the optimal tax and benefit programme. Conversely, the normative analysis makes it explicit that one cannot hope to say how best to design taxes and transfers without both knowing how individuals will respond, and without specifying what one is trying to achieve overall. Often, these two elements are confused in policy debates: right-of-centre policy-makers rarely state explicitly that they have little taste for redistribution, but instead justify their lack of taste for redistribution because they believe that the adverse behavioural responses to high taxes or generous benefits are large. Conversely, left-of-centre policy-makers emphasize the redistributive virtues of benefits and assume that adverse behavioural responses to these and the high tax rates needed to fund them are negligible.

We provide this overview as follows: section 3 introduces the standard optimal tax model developed in Mirrlees (1971). This shows directly how the the optimal tax and benefit system is determined by both the social welfare criterion used by the government and the size of behavioural responses to taxation. Despite the simplifications inherent in the model, we can use it to analyse the optimal tax rate that should apply to top incomes, where we present new, albeit tentative, evidence on the response of top incomes to the large changes in top marginal

tax rates that have taken place in the UK over the last 40 years. Section 4 extends the optimal tax model to allow for labour supply participation effects, and shows that allowing for such responses can drastically change the optimal tax system affecting low income individuals: instead of traditional welfare programmes with high withdrawal rates, large in-work benefits such as Working Tax Credit in the UK or the Earned Income Tax Credit from the US, which can have very low or negative withdrawal rates, can be optimal.⁴ We also discuss the issue of migration and tax design, which can be dealt with in optimal tax models in a similar manner to the issue of labour market participation. Throughout Sections 3 and 4, we make use of the summary of the literature on the behavioural response to taxation provided in Meghir and Phillips, this volume.

In Section 5, we analyze a set of additional issues relevant to the design of the tax and benefit system affecting households. First, we discuss how the family should be taxed: the models considered in sections 3 and 4 abstract from family issues, but a majority of adults in reality live in couples, and so can be assumed to pool income to some extent. We also discuss how the presence of children should be reflected in the optimal tax design. Lastly, we discuss administrative and operational issues concerning benefits.⁵ Section 6 describes how the main elements of the current UK personal tax and benefit system affect incentives to work and earn more and, in Section 7, we provide a critique of the UK tax and benefit system, and set out the direction of reform suggested by the insights from optimal tax theory, and the latest evidence on the behavioural response to taxation. To crystallise ideas, we propose specific changes that could be implemented in the short-run. But most optimal tax theory uses simplified models which leave aside a number of important practical issues such as administrative burden for the government and employers, and ease of use for families.⁶ Those issues have always been important in practice, and the recent “behavioral economics” literature is starting to incorporate them in the analysis. Therefore, we go further and propose a longer-term reform that builds on the short-run changes to incentives by addressing the main practical issues with the current benefits in the UK. Our plan replaces the piece-meal benefits for low-income families (income support, working and child tax credits, housing benefit and

⁴To anticipate our discussion in section 5, the WTC can lead to negative PTRs, but not negative METRs, whereas the EITC can lead to negative METRs as well.

⁵Shaw et al, this volume, discusses administrative and operational issues affecting tax design.

⁶A number of those issues are discussed in more detail in the chapter by Shaw et al, this volume.

council tax benefit) into a single *Integrated Family Support* programme which provides stronger and simpler incentives for work at the bottom, reduces compliance costs for families, and is provided “as-you-earn” and administered in the same way as social contributions through the PAYE withholding system. We show how this can be done in a revenue-neutral fashion, and estimate the behavioural responses to such a reform.

3 The standard optimal income tax model with intensive responses

This section presents the standard model of optimal income tax, based on Mirrlees (1971), in which individuals respond to the tax and benefit system by choosing only how much to work. We then give two applications of the model to the UK:

- first, we can derive an expression for the optimal top marginal tax rate (ie, the marginal tax rate facing the highest income individuals), and we go on to calculate this using new, albeit tentative, evidence on the responsiveness of top incomes in the UK to changes in top marginal tax rates over the last 40 years.
- second, we simulate the entire optimal tax system for the UK given some various highly simplifying assumptions in order to show how the optimal tax system is determined by both the social welfare criterion used by the government, and the size of behavioural responses to taxation.

Before that, though, section 3.1 sets out some of the key terms which will occur throughout this chapter.

3.1 Key concepts

3.1.1 The budget constraint, PTRs and METRs

A useful tool to investigate the disincentive effects of taxes and transfers is the budget constraint.⁷ This shows the relationship between gross earnings (or hours of work) and net income after taxes and transfers, and an example is given in Figure 1A (the example is for a lone parent with 2 children, and we discuss this Figure in more detail and look at other family types in section 6).

⁷This draws on chapter 2 of Adam et al (2006).

The budget constraint contains all the information we need to know about how taxes and transfers affect financial incentives to work, but in this chapter we frequently refer to some summary measures of work incentives:

- the participation tax rate (PTR) is defined as 1 minus the financial gain to work as a proportion of gross earnings. It measures how the tax and benefit system affect the financial gain to work. If someone who did not work had an income from a benefit programme of £60 a week, and would earn £250 in gross earnings, but pay £40 of that in income tax if they were to work, then the PTR is given by $1 - (210-60)/250$, or 40%. The higher the number, the more the tax and benefit system reduces the financial gain to work. A PTR in excess of 1 means the individual would be worse off in work than not working; a PTR equal to 1 means that there is no financial reward to work; a PTR of zero means that the financial reward to work is equal to gross earnings; negative PTRs are possible where benefits are conditional on being in work or having positive earnings.
- the marginal effective tax rate (METR) measures how much of a small rise in gross earnings is lost to payments of tax and reduced entitlements to benefits. It is equal to the slope of the budget constraint at any particular point. The higher the number, the more the tax and benefit system reduces the gain to earning a bit more: a METR in excess of 1 means that an individual would be worse off if they earned a bit more; a METR of 1 means that an individual would be unaffected by any small change in earnings; a METR of zero means that the individual is keeping all of any small rise in earnings; and a negative METR means that an individual's net income increases by more than a small change in earnings (this can arise where benefits act as a proportional subsidy on earnings, such as the phase-in portion of the earned income tax credit in the US).
- It is sometimes more useful to consider the net-of-tax rate, or one minus the METR: this measures how much work pays at the margin.

Figure 1B shows how the schedule of PTRs and METRs for the example budget constraint in Figure 1A; we discuss the particular features of this budget constraint in section 6.

3.1.2 Labour supply responses to taxation

Economists think about the disincentive effects of the tax and benefit system using a labour supply model.⁸ A basic labour supply model assumes that, when deciding whether and how much to work, people trade off the financial reward to working (plus any intrinsic benefits from working) with the loss of leisure time (by "work" we mean "participate in the labour market", rather than doing unpaid work at home or elsewhere).

As we discussed above, taxes and transfers affect labour supply because they alter the financial reward to working, both by making the net wage lower than the gross wage (most taxes, some transfers) and by reducing the financial gain from working compared to not working (most transfers). Economists usually distinguish between two ways that financial considerations affect labour supply:⁹

1. the impact of the METR on labour supply is called the substitution effect, as increasing the METR (thereby reducing the net-of-tax rate) may lead individuals to work less, or to substitute some leisure for work. Economists often measure this effect using the elasticity of earnings with respect to the net-of-tax rate: this measures the percentage increase in earnings following a one percent increase in the net-of-tax rate (see Box 1).
2. In addition, taxes and transfers may also affect labour supply through income effects: higher taxes or cuts in benefits reduce the income available to individuals, and so may induce individuals to work more in order to increase their standard of living. Equally, lower taxes or more generous benefits increase income, and hence may induce individuals to work less. Because the derivation of optimal income tax models is much simpler when there are no income effects (Diamond, 1998 and Saez, 2001), we will assume no income effects in the analysis below, and discuss later informally how the main results change when there are income effects.

⁸See Meghir and Phillips, this volume, and references therein, for more detail.

⁹Meghir and Phillips, this volume, shows the different impacts graphically.

Box 1. The elasticity of earnings

We denote the marginal effective tax rate by τ so that the net-of-tax rate is given by $1 - \tau$. The elasticity of earnings z with respect to the net-of-tax rate $1 - \tau$ is defined as:

$$e = \frac{1 - \tau}{z} \frac{\partial z}{\partial(1 - \tau)}.$$

This elasticity e is always positive. The higher is e , the more responsive are earnings to the net-of-tax rate.

To give an example of its use, if e is 0.2, and the net of tax rate changes from 20% to 25% (ie the METR falls from 80% to 75%), then earnings will rise by $0.2 * \frac{5\%}{20\%} = 5\%$. If the net of tax rate changes from 80% to 75% (ie the METR rise from 20% to 25%), then earnings will fall by $0.2 * \frac{5\%}{80\%} = 1.25\%$

3.2 The Mirrlees model

In the Mirrlees model of optimal taxes, the government is trying to design a tax and benefit system that will maximise social welfare and raise a given amount of revenue. Mirrlees (1971) allowed the tax and benefit system to be non-linear, which means that METRs at a particular point of the earnings distribution can be set to any value without altering METRs at other points. The model assumes that people vary in their earnings potential (or what they would earn if there were no taxes or transfers), and that everyone always works, but chooses how much effort to supply ("effort" can be thought of as hours of work, with a given hourly wage for each individual, but there are other interpretations, as we discuss later).

Before discussing how this model can be used to determine the optimal METR at any point in the income distribution, we first show how it can be used to derive the optimal METR for high-income individuals, a simpler task.

• The optimal top marginal tax rate

To determine the optimal top METR, we will consider the different ways in which a small increase in the top METR affects social welfare. Some of these effects will be positive, and others negative, but at the optimum they must be exactly offsetting, so that no small change in the tax rate can better achieve the goals of the government.

We assume that this top METR applies to earnings above a given level, and we will refer

to this level as the top bracket.¹⁰ There are three impacts on social welfare:

1. With no behavioural response, increasing the top METR will increase government revenue. This is the mechanical effect on tax revenue, and this is a benefit to society, as the revenue can be used for government spending or higher transfers.
2. However, increasing the top METR may also induce top bracket taxpayers to reduce their earnings (but not below the top bracket, because the budget constraint has not changed below this point) because of the substitution effect described above. This is known as the behavioural response on tax revenue, and it is a cost to society as tax revenues will fall.
3. Finally, any increase in the top METR will reduce the welfare of top bracket taxpayers. This is the welfare effect, and it is a loss to society. How large is this loss depends on the redistributive tastes of the government: if the government values redistribution, then, for incomes above a certain level, it will consider that the marginal value of income for top-bracket tax-payers is small relative to the average person in the economy. In the limit, the welfare effect will be negligible relative to the mechanical effect on tax revenue.

An optimal top METR is one where the marginal costs and benefits of increasing it further are balanced. If the welfare effect is negligible, then the government should increase the top METR up to the point where the mechanical increase in tax revenue is equal to the loss in tax revenue from the behavioral response. This effectively amounts to setting the top METR so as to maximize the tax revenue collected from top bracket taxpayers; this can therefore be considered as an upper bound to the top METR above which no government should ever go.¹¹

A precise formula for this optimum top METR is provided in Box 2. The more responsive are earnings to the net-of-tax rate, and the thinner is the income distribution at the top (we formalise this concept in Box 4), then the lower should be the top METR. Later in this section 3.2, we provide estimates for both these parameters for the UK.

¹⁰As already noted, the current top rate of income tax in the UK is 40% and applies to annual earnings greater than £41,435 (in 2008-9). When including National Insurance contributions, the marginal effective tax rate is 47.7% on top earnings.

¹¹It is straightforward to extend the theory to the case where the government has less redistributive tastes and hence the welfare effect is not negligible. See e.g. Saez (2001) .

Box 2. Determining the top rate of income tax

Here we present the optimal marginal tax rate τ for high earners that maximizes tax revenue. We denote by z the *average* income reported by taxpayers in the top bracket (incomes above \bar{z}). By balancing the mechanical and behavioural effects, the optimal rate τ^* can be shown to be given by:

$$\tau^* = \frac{1}{1 + a \cdot e}$$

where a denotes the ratio $z/(z - \bar{z})$ and is a measure of the thinness of the top of the income distribution. The optimal rate is decreasing in both the elasticity e and the shape parameter a . See Appendix for derivation.

• Optimal marginal tax schedule

Using a similar technique to how we derived the optimal METR in the top bracket, we can also derive the optimal METR at any point of the income distribution. As before, the optimal METR at any point is set so as to balance the costs and benefits from changing the METR by a very small amount.

As before, an increase in the METR over a very small band of income has three effects on government tax receipts and welfare:

1. First, the reform increases taxes paid by every taxpayer with incomes above the small band (the mechanical effect).
2. Second, the rise in the METR will reduce earnings for taxpayers in the very small band through the substitution effect, and so generates a loss in tax revenue.
3. Third, the extra taxes paid by every taxpayer with incomes above the small band generates a welfare cost whose size will depend upon the extent to which the government values redistribution.

For an optimal METR, these effects must exactly offset, so that no change in the tax schedule can increase social welfare. An exact expression is presented in Box 3.

The key differences with this analysis and that in the previous section that looks at the optimal top rate are:

- changing the METR at any point affects not just those facing that METR, but also all those with higher earnings

- the welfare cost of extra taxes paid is no longer negligible.

Box 3. Determining the optimal marginal tax schedule

We assume that the government imposes a tax schedule $T(z)$ that depends on earnings z . As shown in Figure 1A and 1B, the slope of this schedule, $T'(z)$, gives the METR when earnings are z . Let $H(z)$ denote the fraction of taxpayers with income less than z (ie cumulative distribution of individuals), and let $h(z)$ denote the density of taxpayers. The optimal tax system is characterized by a grant to those with no earnings (equal to $-T(0)$) combined with a schedule of marginal tax rates $T'(z)$ which define first how the grant should be reduced as earnings increase, and then how additional earnings should be taxed once the grant has been fully tapered away. The government's preferences for redistribution are given by $G(z)$ which measures the social marginal value of consumption for individuals with earnings above z (this should be decreasing in z if the government values redistribution). The optimal marginal tax rate $T'(z)$ is set so as to balance costs and benefits at the margin, and is given by the following formula.

$$\frac{T'(z)}{1 - T'(z)} = \frac{1}{e} \cdot \frac{1 - H(z)}{zh(z)} \cdot (1 - G(z))$$

The optimal tax rate $T'(z)$ is decreasing with the elasticity e , and decreasing in $G(z)$, and increasing in the income distribution ratio $(1 - H(z))/(zh(z))$ which measures the thinness of the earnings distribution. See Appendix for more details.

The formula in Box 3 shows how the optimal METR depends upon the size of the behavioural response to taxation, the government's preferences for redistribution and the underlying shape of the (potential) earnings distribution. In particular, METRs should be higher:

- the less responsive are individuals to the net-of-tax rate
- the more value is placed on redistribution
- at points in the earnings distribution where the number of individuals is small relative to the number of taxpayers with earnings exceeding this amount (this is because the revenue gained from increasing METRs at a given earnings level will be proportional to the number of individuals who have earnings greater than this level; the precise way that we summarise this shape of the income distribution is discussed in Box 4).

Box 4. Summarising the shape of the income distribution

The shape of the income distribution is an important determinant of the optimal structure of METRs. We summarise this shape by the income distribution ratio:

$$\frac{1 - H(z)}{zh(z)}$$

which appeared in the optimal taxation formula presented in Box 3 (where we say that is measured the thinness of the income distribution). The optimal formula shows that the government should apply high marginal tax rates at levels where the density of tax payers, measured by $h(z)$, is low compared to the number of taxpayers with higher income, measured by $1 - H(z)$.

To anticipate the discussion in section 3, it is worth noting that negative METRs are never optimal: if the METR were negative in some range, then increasing it a little bit in that range would raise revenue (and lower the earnings of taxpayers in that range), but the behavioural response (which would be to work less) would also be to raise revenue, because the marginal tax rate is *negative* in that range. Therefore, this small tax reform would unambiguously increase social welfare.

Saez (2001) shows how the analysis changes when income effects are introduced. Income effects encourage work for middle and upper income earners because taxes reduce disposable income, but income effects discourage work for low-income earners, because transfers increase disposable income. Hence income effects make taxing less costly, but make transfers more costly. Therefore, holding other things constant, income effects lead to higher METRs at the upper end, allowing the government to redistribute more, but make redistribution at the low end more costly, and so the net effect on the level of transfers is ambiguous. If income effects are concentrated at the bottom, then they are likely to reduce the size of the optimal transfers at the bottom. If income effects are spread evenly throughout the distribution, then numerical simulations by Saez (2001) show that income effects allow the government to increase the level of transfers paid for by higher METRs across the distribution.

3.3 Empirical evidence on intensive elasticities, and applications to the UK

This section presents two applications of the results shown earlier to the UK tax system. We first derive the optimal top METR, using new, albeit tentative, evidence on the responsiveness of top incomes in the UK to changes in tax rates, based on the response of top incomes

to the large changes in METRs applying to top incomes that have taken place in the UK over the last 40 years. We then derive the entire optimum tax schedule in the standard intensive-responsive Mirrlees model, given assumptions for the labour supply elasticity and the government's preferences for redistribution.

• Top Incomes and the Optimal Top Tax Rate in the UK

Although there is a large literature analyzing the effects of changes in METRs on reported incomes using tax return data in the US (see e.g., Saez, 2004 for a recent survey; some are cited in Meghir and Phillips, this volume), there has been little study of the British case. This is especially surprising, given that the UK experienced a dramatic drop in top METRs. Up to 1978, the top METR on earnings was 83%¹². Under the Thatcher administrations, the top rate dropped to 60% in 1979, and then dropped further to 40% in 1988.¹³

In this section, we propose a very preliminary analysis of the link between top METRs and top incomes, using and extending the top income share series constructed by Atkinson (2007). Those series estimate the share of total personal income accruing to various upper income groups such as the top decile (the top 10 percent), or the top percentile (the top 1 percent), and so they measure how top incomes evolve relative to the average.¹⁴ We have computed the average METR faced by various upper income groups from 1962 to present (in fact, there are two METR series, one including income tax and employer and employee national insurance contributions, and one that also includes the impact of consumption taxes, such as VAT and excise duties.¹⁵

¹²The top rate on capital income was even higher and reached the extraordinary level of 98% from 1974 to 1978, although, as we remarked earlier about the top rate on earned income of 83%, very few individuals had taxable incomes high enough to face this rate.

¹³Dilnot and Kell (1988) try to analyze this issue, but have only access to a single year of micro-tax returns, and rely on aggregate numbers for their time-series analysis. More recently, Blow and Preston (2002) have used micro tax data for 1985 and 1995 to analyze responses to tax rates, but they focus exclusively on the self-employed, and do not look specifically at top incomes. Atkinson and Leigh (2004) have analyzed the link between top income shares and the top statutory marginal tax rate in five English speaking countries including the UK but their study does not estimate effective marginal tax rates and does not focus specifically on the UK case.

¹⁴The definition of income used by Atkinson (2007) (and therefore by us in this section) is close to the broad income definition used in Gruber and Saez (2002), as it excludes capital gains and certain remuneration in kind. However, there are some inconsistencies over time: the most important is that the data represents families before 1990 and individuals after 1990, and we make an adjustment to the pre-1990 data to correct for that (see online appendix for details). Atkinson (2007, p89) also says that the series omits employees' superannuation contributions before 1985, and before 1975-76, the series is net of retirement annuity premiums, alimony and maintenance payments, and allowable interest payments.

¹⁵The consumption tax rate is assumed to be uniform, and estimated using total consumption tax receipts.

Figure 2A displays the METRs (excluding and including consumption taxes) on earnings faced by the top 1% (on the left axis), and the top 1% income share (on the right axis) from 1962 to 2003. It shows an increase in the METR from 1962 to 1978 followed by a dramatic decline in the two key income tax reforms of 1979 and 1988. The top income share series shows an erosion of the top 1% income share up to 1978, followed by sharp upturn starting exactly when the top METR was reduced in 1979, suggesting that top income shares did respond to the lower METR. From a long-term perspective, the top 1% income share doubled from 6% in 1978 to 12.6% in 2003 while the net-of-tax rate (one minus the METR) doubled from $1 - .79 = 21\%$ in 1978 to $1 - .53 = 47\%$ in 2003 (using the rates including consumption taxes). If all the increase in top incomes (relative to the average) can be attributed to the reduction in the METR, this would imply a substantial elasticity almost equal to one.¹⁶

Figure 2B displays the METR and income share of the next 4% (income earners between the 95th and the 99th percentile). In contrast to the top 1%, the METR in 1978 is virtually identical to the current METR; this illustrates that the Thatcher tax reforms cut the progressivity of the income tax within the top 1%, but had relatively small effects on those with slightly lower incomes. However, the income share of the next 4% also shows a break in 1979: the income share is roughly constant at around 12% before 1979, and then increases steadily from 12% to 15% from 1979 to 2003 despite there being little change in the METR.

Two interpretations of this are possible. First, it could be argued that the change in high incomes was not entirely due to the cuts in the METR, and may have been caused by other reforms enacted by the Thatcher administration that were favourable to high incomes. In that case, our previous estimate of 0.93 is biased upward. Second, it is conceivable that income earners in the next 4% group were also motivated to work harder by the prospect of facing much lower rates should they succeed in getting promoted and become part of the top 1% in coming years.¹⁷ In that case, if a cut in the METR facing the top 1% stimulated incomes

These and other computations are described in an online appendix. The METR is an average of the METR on earned and unearned income weighted by the share of earned and unearned income in each group. Our METRs are also weighted by income within each group, as larger incomes have a proportionately larger contribution to the total behavioral response of the income group (indeed, in the optimal top tax rate formula (1), one needs to use the elasticity weighted by income).

¹⁶These elasticities are calculated by computing $\hat{e} = (\log S_1 - \log S_0)/(\log(1 - \tau_1) - \log(1 - \tau_0))$ where S_0 the top 1% income share before the reform, S_1 the share before the reform, τ_0 the marginal tax rate of the top 1% before the reform, and τ_1 the rate after the reform.. In this case, the elasticity is estimated as: $\log(12.6/6.0)/\log((1 - .79)/(1 - .53)) = .93$.

¹⁷Gentry and Hubbard (2004) have tried to estimate such effects in a model of entrepreneurship with US

below the top 1%, our estimate of 0.93 would understate the overall on government revenues.

We show more systematically in Table 1 how this data can be used to estimate the elasticity of broad income with respect to the net-of-tax rate. The first 2 rows of Table 1 focus on the 2 key tax cuts of 1979 and 1988, and compare 1978 with 1981 and 1986 with 1989, respectively¹⁸. Column (1) estimates the elasticity of top 1% incomes by calculating how the share of income received by the richest 1% of individuals changes relative to the change in the METR that this group was subject to. It shows positive, but not very large, elasticities of 0.34 and 0.26. However, as we discussed above, the longer-run perspective suggests higher elasticities. Indeed, the third and fourth rows compare years 1962 to 1978 (when METRs for the top 1% increased) and years 1978 to 2003 (as we discussed above), and these comparisons imply substantially higher elasticities of 0.61 and 0.93. Finally, the bottom row presents the coefficient of a simple time-series regression of the income share of the top 1% on the METR. Rather than just comparing the changes between two different years, this approach uses data over the entire 1978 to 2003 period, and suggests an elasticity of 0.73 (which is statistically significant). In Column (2) we again calculate the elasticity estimates of top earners, but we exclude consumption taxes from our measure of METR: this hardly changes the elasticity estimates (because average consumption tax rates have changed by much less than the marginal rate of income tax applying to top incomes).

The elasticities reported in columns (1) and (2) are unbiased estimates only if, absent the tax change, the top 1% income share would have remained constant. As we explained above, this assumption seems contradicted by the fact that the top 5-1% income share increased from 1978 to 2003 in spite of no change in METRs. If we assume that, absent the tax change, the top 1% share would have increased as much as the top 5-1% share, we can calculate what is referred to as a *difference-in-differences* estimate, which is presented in column (3) of the table.¹⁹ These difference-in-differences (DiD) estimates are smaller for the long-term 1978-2003 comparison, and for the full time-series regression, although they remain substantial at 0.64 and 0.46 respectively. It is conceivable that, absent the tax change, the top 1% share would still have increased more than the top 5-1% share, perhaps because the Thatcher administration

data.

¹⁸We do not use 1990 because of the change from couple to individual tax filing which creates a small discontinuity in the Atkinson series.

¹⁹Those elasticity estimates are $\hat{e} = (\log S_1/S_1^c - \log S_0/S_0^c) / (\log(1 - \tau_1)/(1 - \tau_1^c) - \log(1 - \tau_0)/(1 - \tau_0^c))$ where S^c and τ^c are the income share and marginal tax rate for the “control group” Top 5-1%.

implemented other policy changes favorable to top incomes, or because of structural changes in the labour market and changes in the returns to human capital.

Table 1: Elasticity Estimates for Top Income Earners

	Simple Difference (1)	Simple Difference (excluding consumption tax from MTR) (2)	DD using Top 5-1% as control (3)
1978 vs. 1981	0.34	0.32	0.08
1986 vs. 1989	0.37	0.38	0.41
1978 vs. 1962	0.61	0.63	0.86
2003 vs. 1978	0.93	0.89	0.64
Full time-series regression (s.e. in brackets)	0.73 (0.13)	0.69 (0.12)	0.46 (0.13)

The second parameter in formula (1) is a , the measure of the thinness of the income distribution at the top (see Box 4). Figure 3 shows how our measure of the shape of the income distribution (discussed in Box 4 above) varies with earnings in the UK: the hazard ratio is very high at the bottom, falls as income increases, and then rises slightly until it becomes flat around 0.6, implying a value of a of $1/0.6 = 1.67$.

What do these estimates mean for the optimal top rate in the UK? We gave an expression for the optimal top rate in Box 2 of $\tau^* = \frac{1}{1+a-e}$. With $a = 1.67$ and an estimate of $e = 0.46$, the revenue-maximising top rate is 56.6%, only a little higher than the the actual total top METR in 2008/9 (around 53% including consumption taxes).²⁰ But we would stress that, as our estimate of the elasticity is tentative, so is the estimated optimal top rate. Taking values of the elasticity 1 standard deviation either side of the central estimate gives a range for the optimal top rate of 50.4% to 64.5%. But our analysis is also consistent with the current top METR being too high: using the value of the elasticity from the simple difference over the period 1978 - 2003 would give an optimal top rate of 40.2%, and using the difference-in-difference estimate of the elasticity from the same period would imply an optimal top rate of 49.4%, slightly lower than the actual top rate. Indeed, both these estimates imply that cuts in

²⁰The revenue-maximising top rate is the optimal top rate if the government places no cost on the top 1% having less income as a result of the tax rise.

the METR facing the richest 1% in the UK would actually increase tax revenues (although see Box 5 for a discussion of the difference between taxable income and broad income elasticities).²¹

Box 5. Taxable income and broad income

Note that to estimate the revenue implications of raising the METR that applies to the top 1% of earners in the UK given all other aspects of the current UK tax regime, one would want to use a taxable income elasticity (which measures how income that is subject to income tax changes when the net of tax rate changes). But the income measure used in our analysis was close to a broad income measure, rather than taxable income (so it includes some sources of income not subject to income tax). For optimal tax design, the the right concept to use is a broad income elasticity, because the difference between broad income and taxable income is a function of the tax system and enforcement efforts, and therefore depends entirely on the choices made by governments. For the same reason, the taxable income elasticity is unlikely to be constant across income tax regimes. For example, we might expect the taxable income elasticity to be higher in the US, than in the UK, because there are more opportunities to reduce taxable income in the US tax code than in the UK. In the UK, the main ways in which one can reduce taxable income would be through higher contributions to private pensions (which to some extent represent only deferred taxation because eventual pension income is taxable), and through charitable giving (to which there may be externalities).

This first-pass analysis shows that identifying the elasticity of top incomes, a key ingredient in the optimal tax rate formulas derived above, is not simple. The evidence is consistent with significant behavioral responses to top taxpayers to METRs, certainly suggesting that the key elasticity is not zero. As the formula (1) shows that the upper bound on METRs depends critically on the level of this elasticity, it would be very valuable to explore this issue in more detail using the rich UK tax return micro-data (the Survey of Personal Incomes) that have now become available to researchers. Unfortunately, there has been no large change in METRs since 1988; and without such a change it is extremely difficult to estimate this elasticity. It is conceivable that these behavioural responses change over time; see, for example, the discussion of migration effects below.

Note also that these calculations have only derived the optimal rate for the richest 1% of the population. For many years, the highest rate of income tax in the UK has applied to a much greater proportion: in 1991/2, 3.5% of adults paid income tax at the highest rate, and

²¹In 2004/5, the richest 1% of adults, or 470,000 individuals, had incomes in excess of £100,000, with a mean of £156,000: see Brewer et al (2008).

this has risen to 6.8% in 2004/5, and almost 8% by 2007/8.²² This means that the conclusions in this section should not be seen as implying that the existing higher-rate of income tax with its existing thresholds should be changed: as the section below shows, the optimal METR that applies to, say, people in the top 6% of income earners but outside the top 1% could be lower or higher than the optimal METR at the top.

• Simulations of the whole optimal tax system in the UK

Having estimated the optimal top METR, we below simulate the whole optimal tax structure using the Mirrlees model set out in the previous section, and based on the actual UK earnings distribution, and various assumptions about the intensive labour supply elasticity (full details are in the Appendix). The simulations attempt to show the optimal tax schedule which provides total net tax revenues equal to the current tax system, including revenue from individual income tax, NICs, and consumption taxes, net of spending on existing transfers for families with children or those with disabilities.²³ To focus specifically on income tax, we have computed the optimal income tax schedule when we keep consumption taxes (VAT and excise taxes) at their current level (around 17% on average), which we assume to be constant as income varies. The simulations assume that the tax and benefit system is at an individual level.

Figure 4A shows the optimal income tax schedule, exclusive of consumption tax, assuming a constant elasticity of 0.25 and with the government valuing redistribution (we define this more precisely in Box 6). It shows that for very low levels of earnings, individuals face a METR of around 70%; the METR then decreases relatively quickly with income, reaching 36% as incomes approach £30,000 per year. As incomes increase further, so too does the METR, eventually settling at around 64% for incomes above £200,000.²⁴

The U-shape pattern of optimal marginal tax rates is not surprising in light of our theoretical discussion: it is driven by the U-shape of the hazard ratio $(1 - H)/(zh)$ (see Box 4; this describes the thinness of the income distribution), as well as the decreasing shape for $1 - G(z)$, the government's preferences for redistribution, both combined with the assumption that the

²²See HMRC (2007c).

²³We assume total transfers are equal to the amount spent on Job Seekers Allowance, income tax credits and reliefs, child benefit, housing benefit, council tax benefit, and income support.

²⁴These marginal rates will increase once we consider the impact of consumption taxation: for example, consumption taxes act to effectively raise the marginal rate of high incomes from 64% to 70%.

elasticity does not vary with earnings.

We now consider how our views regarding the optimal schedule depend on the labour supply elasticity. Meghir and Phillips (this volume) survey the elasticity of hours worked with respect to the wage. For men, they say that “although one can start discussing the relative merits of the approaches taken, existing research will lead to the conclusion that the wage elasticity is zero.” For women, they conclude that the elasticity of weekly hours worked is “in the range of approximately 0.0 to 0.3”. Their preferred estimate is a value of 0.13 for all married women except those with young children (for those with children aged 3-4, the value is 0.37). They also say that “the results of annual labour supply show greater responsiveness to wages”, probably because variations in annual hours worked are a combination of participation responses (whether a woman works at all in a given week), and intensive responses (changes in the hours worked per week).

But hours worked are not the only way in which taxable income can respond to tax changes. For many individuals, the idea that the hourly wage cannot be affected by the amount of effort expended by the individual (as assumed in the theoretical models in Section 3) is too simplistic; earnings could respond to tax changes through changes in the hourly wage (whether through bonuses, tips, job changes or even by workers on piece rates working faster) as well as hours worked. Taxable income reported to the revenue authorities, though, is not the same as gross earnings, and can vary in response to tax changes through changes in the form of compensation, the response of non-labour income, and changes in the amount of income reported to the tax authorities, whether through avoidance or evasion. Saez (2002) argues that “elasticities of earnings with respect to the tax rate [at the bottom end] are ... perhaps around 0.25”, and that: “there is little consensus about the magnitude of intensive elasticities of earnings for middle income earners, although this elasticity is likely to be of modest size for middle income earners and higher for high income earners. Gruber and Saez [2002] summarize this literature and display empirical estimates between 0.25 and 0.5 for middle and high income earners.” (p. 1057, Saez (2002)), although most of this is focused on the US.

Figure 4A also displays an optimal schedule in the case where individual labour supply is more responsive to changes in income (an elasticity of 0.5). The figure demonstrates that this would produce lower METRs across the earnings distribution, falling as low as 20%, with a top rate of 45% (slightly below the existing rate). The intuition for the difference here is

simple: when individuals are more responsive to tax changes, they will react more to a given METRs (reducing their labour supply by more), and this places a limit on how high METRs can go. Correspondingly, and as shown in Table 2, the benefit programme is less generous when the elasticity is higher.

Finally, we consider how the government’s preferences for redistribution affect the optimal schedule (see Box 6 and the online appendix for more detail). An interesting case to consider is known as the Rawlsian case, which seeks to maximise the welfare of the least well-off member of society.²⁵ As Figure 4B and Table 3 show, under this criteria, we would have a higher lump-sum grant and higher METRs across the entire distribution of earnings. Hence, rates are higher at the bottom, and are the same as the utilitarian case at the top. Therefore, with a Rawlsian criterion, the optimal shape becomes closer to an L than U-shape.

Box 6. Expressing the preference for redistribution

In calculating social welfare, we first transform (money metric) utilities so that we allow for the possibility that the government attaches more weight to the welfare gains of individuals whose level of utility is initially low. A convenient and simple way of capturing this concern for inequality is to transform original utilities u as follows:

$$\begin{aligned} & \frac{u^{1-\gamma}}{1-\gamma} && \text{if } \gamma \neq 1 \\ & \log(u) && \text{if } \gamma = 1 \end{aligned}$$

Social welfare is then obtained by summing these transformed utilities across individuals. Whenever γ is positive, any increase in utility translates into a less than proportional increase in social welfare. When $\gamma = 1$, which is the case that we consider here, the government is placing twice as much weight on the utility gains of an individual relative to another individual whose utility is twice as high. If concerns for inequality were even stronger, represented by say $\gamma = 2$, then they would be placing four times as much weight on the utility gains of the less well off individual. When $\gamma = 0$, there is no concern for inequality; when γ gets very large, only the worst-off individual in society determines social welfare (the Rawlsian case discussed further below).

The form of individuals’ utility function is given in the online appendix, but note that it is quasi-linear in income, and so it does not display diminishing marginal utility of income, which can by itself provide a motive for redistribution even if a government has a strictly utilitarian social welfare function.

²⁵The Rawlsian criterion can therefore be seen as a bound on the maximum level of redistribution that the government wishes to do; note that in the optimal tax model, even the Rawlsian government has to raise revenue for some reason, and this places a limit on the size of the transfer to the poorest in society.

Table 3: Optimal Tax Rates and Lump-sum grants

Redistribution strength	Elasticity	Average MTR	Lump-sum grant (per year)
$\gamma = 1$	0.25	45%	£5,580
Rawlsian	0.25	73%	£8,150
$\gamma = 1$	0.50	31%	£4,270
Rawlsian	0.50	58%	£6,760

4 Optimal taxes and transfers when there are participation effects

The model described in the previous section assumes that individuals respond to the tax and benefit system only by varying their earnings as a function of the net-of-tax rate they face (known as the intensive margin/response). However, changes in whether people participate in the labour market at all (known as participation or extensive responses) are poorly captured within such a framework (see Blundell and MaCurdy, 1999). Indeed, following a small increase in the net gain of work, people tend to enter employment at, say, twenty or forty hours a week, rather than one or two hours. Such extensive labour supply responses are particularly important at the bottom of the income distribution, and can be incorporated into a model of labour supply using fixed costs of work (Heim and Meyer, 2004).

Participation effects are important: accounting for them radically modifies the structure of optimal taxes for low income families from the one obtained above (Diamond (1980), Saez (2002)). In this section, we outline the key theoretical results, and then discuss recent applications using UK data.

4.1 Theory

We continue to work with a simple labour supply model, but this time individuals only choose whether or not to work, and this decision depends on the relative rewards to working and not working (including the costs of work). The responsiveness of this decision to the financial rewards can be summarised in the elasticity of participation with respect to the net return to work, similar to the elasticity defined in Box 1. If individuals do work, their earnings are fixed.

If the government implements a tax and benefit schedule that determines the disposable income of individuals both in and out of work, then an individual chooses to work if the net

return to work exceeds her cost of working.²⁶ Consider the impact of a small rise in the PTR at a given level of earnings. This reform affects only individuals with this earnings potential, because there are no intensive responses. As in section 3.1, this reform has three effects on government tax receipts and welfare:

1. the reform increases the taxes paid by every taxpayer at the given level of earnings who works, increasing government revenues.
2. those extra taxes reduce the welfare of the workers who pay this extra tax, with the value that the government places upon this dependant upon its redistributive preferences.
3. the tax rise induces some of the workers at this earnings level to drop out of work, and this has a cost.

At the optimal, these effects must balance. In the Appendix, we derive formally what this means for the optimal tax schedule, with the result presented in Box 7.

Box 7. Optimal tax rates with participation responses

Let $g(z)$ denote the value the government places on increasing income of individuals with income z . If the government values redistribution, then $g(z)$ will be fall as z rises.

Defining the participation tax rate as $t(z) = (T(z) - T(0))/z$, we can derive the optimal tax rate as:

$$\frac{t(z)}{1 - t(z)} = \frac{1}{\eta} \cdot (1 - g(z))$$

This formula is a simple inverse elasticity tax rule for the participation tax rate on work. The PTR decreases with the elasticity η and with $g(z)$, the social value of marginal consumption for individuals earning z .

As Box 7 shows, the optimal average tax rate at any given earnings level will be lower:

1. the more highly the government values income at that earnings level
2. the higher is the participation elasticity (since it is not desirable to tax individuals who adversely respond to reductions in their incomes)

²⁶Since individuals of a given ability level may differ in their costs of working, for any given tax system some of these individuals may choose to work and others not.

A striking implication is that, if the government values redistribution - so that $1 - g(z)$ is negative - then the participation tax rate should be *negative* for low earnings - in other words, low income workers should receive an earnings subsidy. Hence, in sharp contrast to the intensive model, the extensive model implies that earnings subsidies or work-contingent credits (such as the earned income tax credit or the working tax credit) should be part of an optimal tax system.²⁷

The intuition for this result can be understood as follows. Starting from a tax and benefit system with a positive participation tax rate for low skilled workers, and suppose the government contemplates strengthening work incentives for low skilled individuals by reducing the taxes that they would pay when working. This has the following effects:

- the cut in taxes means that tax revenues fall, and this is a cost
- but the associated increase in income of low skilled workers is viewed positively by a government that values redistribution, because it would prefer these individuals to have income more than the average individual. By our assumption that we are considering low income individuals (for whom $g(z)$ exceeds 1), this benefit effect has to outweigh the costs of reduced revenue.
- the behavioural response from cutting the PTR is to induce some low-skilled individuals to start working, and this increases government revenue (because the individuals who move into work pay positive net taxes).

Hence, this reform is unambiguously desirable, and the implication is that positive PTRs for low-income workers cannot be optimal.

These arguments were true for a model where the only response is along the extensive margin. A more realistic model which allows for both intensive and extensive effects is presented in Saez (2002). To summarise the implications of such a model, consider the situation outlined above, where the government lowers taxes (which are currently greater than zero) for low-skilled workers. If there are both intensive responses and extensive responses, then cutting taxes here would induce some higher skilled workers to reduce their labour supply,

²⁷This result is robust to introducing income effects, as formula (3) remains valid with income effects: see Saez, 2002.

as well as inducing some non-workers to work. Although the latter response is a benefit to society, the former is a cost, and so cutting taxes has ambiguous effects on labour supply and therefore overall social welfare. A government contemplating strengthening incentives by cutting taxes facing low-income workers must therefore weigh precisely the positive participation effect and the negative intensive labour supply effect, and the model in Saez (2002) gives a precise formula for that trade-off.

- **Interpreting the participation response**

The extension of the optimal tax model to allow for non-participation has other applications. Two of those are tax evasion and migration.

To apply the model to tax evasion, our earlier concept of "earnings" could be interpreted as "earnings reported to the government agency administering taxes and transfers". Suppose that low income earners can decide to either work in the formal sector, where we assume full compliance with the tax and benefit rules, or in the informal sector, where we assume full non-compliance. In that case, the decision to work or not work can be replaced by the decision to work and report earnings, or to work informally and not report earnings. In that case, for a given level of tax enforcement efforts, our earlier analysis (and the formula presented in Box 7) remains valid. However, in such a model, the government might recognize that some of all individuals reporting no earnings are in reality working informally, and so might actually be better off than low-income workers in the formal sector. This may lead the government to place a lower value on the consumption of individuals with no reported earnings than they do on workers with low reported earnings (ie, $g(0) < g(z)$ for some z), and this would make subsidies for work even more likely to be optimal.²⁸

Second, taxes and transfers might affect migration in or out of the country. For example, high tax rates on skilled workers in continental Europe might induce some of them to move to the UK or the US where the burden of tax on high-income individuals may be lower, and generous benefits for lower income individuals in certain countries might encourage migration

²⁸However, subsidies for low income individuals might induce individuals to over-report self-employment income. In the US, Saez (2002b) shows that the self-employed are much more likely than wage earners to report income which makes them entitled to the maximum EITC payments, strongly suggesting that self-employed individuals manipulate their reported earnings to take advantage of the EITC, making use of a flexibility not available to wage earners.

of low skilled workers toward those countries.²⁹ In the online appendix, we discuss how the migration decision can be incorporated into optimal taxation models.

4.2 Evidence on extensive elasticities and empirical applications

Meghir and Phillips (this volume) show that there is a wide range of participation elasticities for women in the literature: “Aaberge et al. (1999), Arrufat and Zabalza (1986) and Pencavel (1986) find results of 0.65, 1.41 and 0.77-0.89 respectively using cross-sectional data-sets from Italy, the UK and the US, and using significantly different modelling and estimation strategies....Devereux (2004), however, finds a lower degree of responsiveness with the elasticity at the median family income equal to 0.17.” There is consensus, though, that participation is more elastic amongst women from poorer families so that “participation is likely to be the key margin of adjustment for poorer women”.

For men, the two studies of static labour supply cited by Meghir and Phillips (this volume) suggest an elasticity close to zero, but they highlight that a separate literature on the effect of unemployment benefits on the duration of unemployment has consistently found that higher benefits lead to a longer period out of work. Even including these effects, though, they suggest the overall participation elasticity for men is very close to zero, at 0.04. A dynamic model for young men in Germany (Adda et al. (2006)) gives a similarly low participation elasticity (0.06). But Meghir and Phillips also provide their own, new empirical evidence. This very clearly shows the heterogeneity of responses: for highly educated men, it is hard to reject the idea that the participation elasticity is zero, but the estimate for men with low educational qualifications is 0.23 for single men, and 0.43 for men in couples.

There are very few empirical studies of optimal tax systems that incorporate intensive and extensive responses: Saez (2002) is an example for the US. Of course, one approach to the second goal of this chapter - where we seek to apply the lessons from optimal tax theory to make recommendations for the UK tax and benefit system - would have been to use an optimal tax model that allowed for intensive and extensive responses to solve for the optimal schedule. We have not taken this approach, though, primarily because we needed to reflect that the current tax system in the UK has different tax schedules for single people and couples, and

²⁹Clearly, governments can use other tools to affect immigration, and such policies are taken here as given. In the EU, emigration and immigration across EU countries is almost completely deregulated, and so our analysis is particularly relevant in this context.

schedules that vary by the number of children, but also because we also consider the impact of the tax and benefit system on family formation and fertility, and administrative issues, and it is to these we turn in the next section.

5 How should taxes and benefits treat "the family"?

The models we have considered thus far were based on individuals and so abstracted from family issues. In reality, a majority of adults live in couples, and can be assumed to share income to some extent. In this section, we discuss how the family should be taxed, and how the presence of children should affect taxes and benefits. See Boxes 8 and 9 for a discussion of how these issues are currently treated in the UK.

Under a pure individual-based taxation, tax liability is assessed separately for each family member and is therefore independent of the presence or income of other individuals living in the family or household. At the other extreme, in a system of fully joint taxation of couples, tax liability is assessed at the family level, and depends on total family income (this is how income tax works in the US, for example). Over the past three decades, there has been an international trend from joint to individual taxation of husbands and wives, and today the majority of OECD countries use the individual as the basic unit of taxation (income tax in the UK moved from being family-based to individual-based in April 1990). But tax credits and transfers for low-income families in the UK are based on total family income, as are the equivalent welfare benefits in most other OECD countries, and there has been much less impetus to move to an individually-based system for assessing transfers. Of course, there are many other ways of designing a tax and benefits schedule for individuals that varies with the presence of a partner than a fully joint tax and benefit system, and many EU countries with individual tax systems have some form of recognition of marriage or the presence of a partner (see Di Tommaso et al. 1999).

In general, there are several important points to be considered when designing taxes and benefits for individuals who can live either alone or in a couple:

1. if there is any sharing of resources within a family, a person with a low income living with a high-income spouse is better off than an otherwise-equivalent person living with a low-income spouse. Therefore, if the government values redistribution, two adults earning the

same ought not to be taxed the same if their partners' incomes are very different. This redistributive principle is achieved to a limited extent by having a progressive income tax system based on family income, since it imposes higher average tax rates on adults with high-income spouses than on otherwise-identical people with low-income spouses. By contrast, an individual-based income tax does not meet this redistributive criterion: it imposes the same tax burden on individuals irrespective of their partner's earnings.

2. if there are economies of scale in households, so that two adults living apart could achieve the same standard of living with less income if they lived together, then this arguably provides a reason for the tax system to take account of whether individuals are in a couple or not, presumably by taxing individuals more when they are living as a couple than when they are living alone.
3. family-based income tax and benefit systems are highly likely to create a marriage (or cohabitation) subsidy or penalty, as the net tax liability of the two adults might change if they decide to cohabit or marry. This is well-documented in the US, where the income tax system is family-based for married couples, but not for cohabiting couples. Because the US tax system is progressive - in other words, the tax burden rises as income rises - couples with very unequal incomes (such as single-earner couples) benefit from a marriage subsidy, while couples with similar incomes (such as two earner couples) face a marriage penalty. Although the marriage penalty/subsidy attracts substantial public attention, it becomes relevant for optimal taxation only if the decision to marry is sensitive to those fiscal incentives. Hoynes (this volume) concludes that "overall, the research [mostly using US data] finds tax effects on marriage that are consistent with the theoretical predictions but are small in size" (these studies are cited in Eissa and Hoynes (2004)); a related literature finds that marriage is also sensitive to the financial incentives inherent in welfare systems (almost always anti-marriage incentives), but that the elasticities are small (see Hoynes (1997)). Overall, then, Hoynes concludes that "the estimated elasticities with respect to the tax-induced financial incentives to marry (and divorce) are small". However, even if marriage or partnership decisions are relatively insensitive to fiscal consequences, we might expect that how individuals report their family circumstances to the government authorities would be affected by sufficiently

large cohabitation penalties or subsidies.³⁰

4. the empirical literature has shown that the labour supply of secondary earners is more responsive to taxes than that of primary earners (see Meghir and Phillips, this volume, or Blundell and MaCurdy, 1999). Therefore, the earnings of secondary earners should be taxed at a lower rate than the earnings of primary earners for efficiency reasons.³¹ This goal is achieved to some extent by a progressive individual-based income tax, since primary earners have higher incomes and hence tend to face higher METRs than secondary earners. By contrast, a family-based tax and benefit system generates identical METRs across members of the same family, and thus does not meet this efficiency principle.
5. Any tax and benefit system other than a fully joint system will give individuals in a couple a tax incentive to equalise their asset holdings or income streams. While this might be a deliberate policy intention, a consequence is that it gives couples an opportunity to reduce their tax burden by transferring assets from an adult with the higher METR to the lower METR.

Deriving general optimal tax results for couples is, in general, extremely complicated. Kleven et al (2006) consider a simple optimal tax model for couples where the primary earner chooses only how much to work (as in the models outlined in Section 2) and the secondary earner chooses only whether or not to work (as in the models outlined in Section 3). In contrast to the separable and linear tax system in Boskin and Sheshinski, 1983, they consider a fully general joint taxation system. Naive intuition based suggests that, for redistributive reasons, the participation tax rate on the secondary earner should be higher when the earnings of the primary earner are larger, as the contribution of the secondary earner's income to the family's well-being is minimal. However, the authors show that the reverse is true: the participation

³⁰In the UK, Her Majesty's Revenue and Customs and the Department for Work and Pensions both estimate the extent of money lost to such fraud or error relating to the presence of a partner: these estimate that £67m was overpaid in income support, jobseekers allowance and the pension credit, £30m in housing benefit (both in in 2005/6) and £320m overpaid in tax credits (in 2004/5) (DWP, 2007a and HMRC, 2007b). Powerful circumstantial evidence that such fraud exists comes from the fact that the UK government is paying child-contingent support to between 5 and 10 per cent more lone parent families than are thought to live in the UK (Brewer et al, forthcoming; Brewer and Shaw, 2006).

³¹This is in line with the traditional Ramsey principle of optimal taxation that commodities with relatively more elastic demands should have relatively lower tax rates. See also Boskin and Sheshinski, 1983, Alesina and Ichino, 2007.

tax rate on the secondary earner should be decreasing with the earnings of the primary earner and, symmetrically, the primary earner should face a lower METR if his spouse works.

The correct intuition is the following: conditional on the earnings of the primary earner, two-earner couples are always better off than one-earner couples. Hence, the government would like to redistribute from two-earner couples to one-earner couples. The value of such redistribution is larger for couples with low primary earnings because the contribution of the secondary earner to household utility is then more important. Therefore, the redistributive virtue of taxing secondary earnings is actually higher at the bottom of the primary earnings distribution, explaining why the tax rate on secondary earner is decreasing with the primary earner income. If the tax schedule for two-earner couples is seen as the base schedule, the optimal schedule for one-earner couples is obtained from that base schedule by introducing a tax allowance for non-working spouses that is larger for couples with low primary earnings than for couples with high primary earnings. This shrinking tax allowance generates an implicit tax on secondary earners which decreases with primary earnings.

This result suggest that a progressive joint income tax system goes in the wrong direction, and that neutral individual taxation is closer to the optimum. However, it is important to note that, in practice, benefits for low-income families are almost always based on joint family income, and the phasing-out of those programmes creates implicit taxes on secondary earners which are decreasing with primary earnings. For example, a secondary earner in the UK with modest earnings would face a relatively high (average and marginal) tax rate when her partner's earnings are low, because the second adult's earnings reduce the family's tax credit entitlement as well as being subject to income tax and national insurance contributions, and would face a relatively low (average and marginal) tax if her partner's earnings are high, because the secondary earnings are subject only to the individual income tax and NICs. Hence, the results in Kleven et al. suggest that the broad way in which tax and benefit systems of many OECD countries treat the incomes of a couple, including that of the UK, are consistent with optimal tax results.

Box 8. Marriage and cohabitation penalties in the UK tax and transfer system.

As we discuss in chapter 6, the UK has individual assessment for income tax, but welfare programmes and the child and working tax credit are assessed against the joint income of co-resident couples, where legally married or not. There are currently very few tax-induced marriage penalties or subsidies in the UK.^a However, there are substantial so-called cohabitation or couple penalties or subsidies in the UK: these arise because welfare benefits and tax credits are assessed against the joint income of cohabiting couples, whether legally married or not.^b The same structural features that lead to such cohabitation penalties also give differences between the incentives to work facing the first and second earner in a couple. Typically, the PTR of the second earner will be considerably lower than that of the first earner (see also Figure 6A), but the direction of recent reforms has tended to increase PTRs of second earners as entitlement to tax credits has risen in real terms (see Brewer, 2007).

^aSee Bowler (2007); some groups have argued that there should be more subsidies: see Social Justice Policy Group (2007).

^bSuch couple penalties or subsidies are usually shown either by calculating the change in net transfers from the state that two adults would experience if they were to cohabit - a complicated calculation that requires assumptions on how housing costs and labour supply would change upon cohabitation - or by calculating the change in net transfers that a cohabiting couple would experience if they (fraudently) claimed to be living apart. See Anderberg et al. (2007); see also Kirby (2005) and Draper (2007)

• **Collective labour supply model**

How disposable income is allocated among family members raises interesting issues. Empirical findings by Lundberg et al. (1997) show that giving an allowance for children directly to the mother instead of giving it to the main earner through a reduction in taxes increased spending on children significantly. This shows that families do not fit what is called the "unitary model", whereby a family acts as if all the adults in it care about the same things. Many other models of behaviour of couples are possible, and Chiappori (1988, 1992) developed a model where consumption is allocated within family members in an efficient way (so that it is not possible to make one member in the family better off without making another worse off), but that the power each family member has in the decision-making process depend on their relative incomes or on whom is entitled to the government transfers (this is known as a collective labour supply model).

How does this affect our analysis? Suppose, for example, that husbands have too much power within a couple, and get too much control over how income is used than their spouses,

