
The Taxation of Labour Earnings

The taxation of earned income accounts for more revenue than any other form of taxation across most OECD countries. Alongside the benefit and tax credit system, it also does most of the heavy lifting in redistributing resources from richer to poorer households. Not surprisingly, it occupies a special place in debates about levels and structures of taxation.

In this and the next two chapters, we look in detail at earnings taxation, and relevant aspects of the welfare benefit system which are intended to support those who are out of work or earning little. In this chapter, we provide some relevant background on earnings inequality and labour force participation. We go on to explain the inevitable trade-offs between redistribution to those with low incomes and/or high needs and the need to promote economic efficiency, especially work incentives. We consider how to think about those trade-offs in an optimal tax framework and in the context of the family, not just the individual. The key role that the limited information available to the government plays in restricting policy options is emphasized. In the subsequent chapters, we look explicitly at reform options for the UK.

In doing that, we need to bear in mind throughout that it is social welfare we are interested in. That depends on much more than measured income. There is a cost to people in working more hours. Needs vary and income may be more valuable to some than to others. Governments will tend to value income being transferred from rich to poor and society may care about

how (un)equal is the well-being of different people.¹ Measures of social welfare tend to acknowledge that society generally places more weight on gains to the poor and needy.

The maximization of this measure of 'social welfare' stands a long way from the maximization of some crude measure of aggregate income. It allows for work effort, needs, and inequality. However, by acknowledging that people may respond to tax changes by adjusting the level of their work effort, it also explicitly acknowledges the importance of incentives to generate income in the economy. A balance has to be struck between the level of work effort generating earnings in the economy—the size of the pie—and the degree of redistribution to the poor and needy.

If society cares about inequality and poverty, there will always be some willingness to sacrifice a part of national income in order to achieve distributional objectives. The fundamental design issue is to minimize such losses while raising sufficient revenue to finance desired public services and satisfy concerns over inequality and poverty.

Our goal in this chapter is to lay out principles for balancing economic efficiency and redistributive objectives in a coherent tax and benefit system. To do this, we will need to address two key questions. The first is positive: how do individuals' earning decisions respond to taxes and benefits? The second is normative: given how people behave, what tax and benefit system would best meet policy goals?

In choosing a set of tax rates, a government may not be concerned only with redistribution and economic efficiency. It may also make social judgements—for example, wishing to favour the working poor over the non-working poor, or married parents over cohabiting ones. Even so, some tax arrangements dominate others by achieving the particular social objectives in a more efficient manner.

Over the longer term, the taxation of earnings can also affect choices over careers, education, and training and over whether to be self-employed or an employee. All affect the pattern of earnings and employment over the working life of an individual. Consequently, in our design of the rate

¹ The objective of maximizing the aggregate of all individuals' well-being is the form of utilitarianism proposed by Bentham (1789). At the other extreme, the view that so much weight should be given to equality that the quality of a society should be judged only by the well-being of its worst-off member is often attributed to Rawls.

schedule, we have to balance a broad set of incentive effects with the goal of raising tax revenue and redistributing to those with greater needs. By looking at the tax system as a whole, the aim of this book is to present recommendations for a tax system that provides a coherent treatment of savings, pensions, human capital investments, and earnings.

3.1. EMPLOYMENT AND EARNINGS

In this section, we describe very briefly two important underlying trends which will inform much of the rest of what we say about earnings taxation—changes in employment rates and changes in the distribution of earnings.

Employment rates have changed a lot over the last three decades. Figure 3.1a shows this for men in the UK. There is a systematic fall in employment at nearly all ages over this period. Even though employment rates edged up in the last decade before the recent financial crisis, they did not come close to

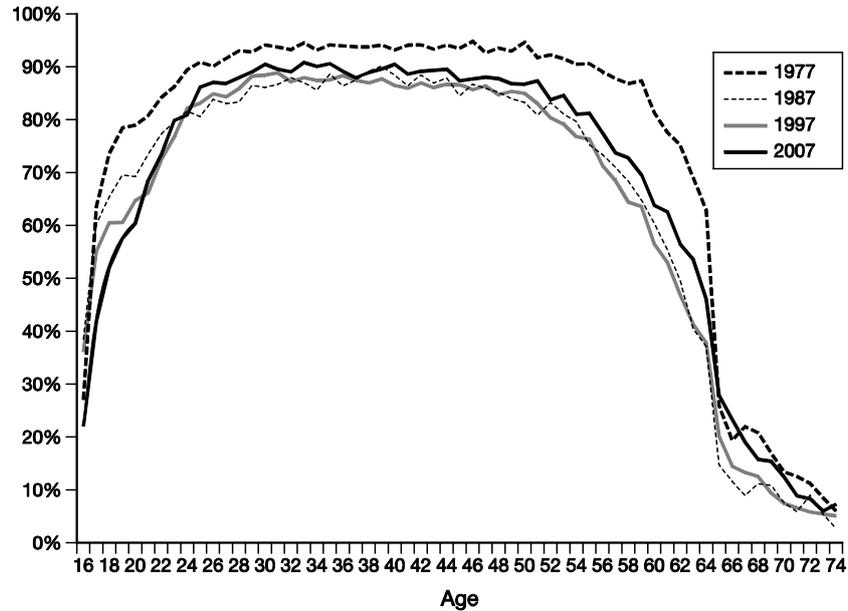


Figure 3.1a. Employment rates by age in the UK over time: men

Source: Blundell, Bozio, and Laroque, 2011.

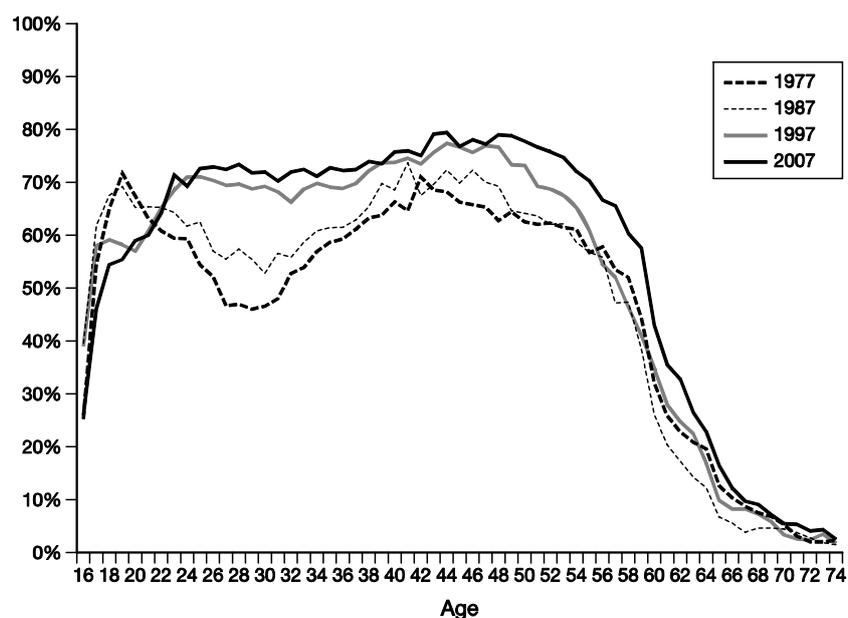


Figure 3.1b. Employment rates by age in the UK over time: women

Source: Blundell, Bozio, and Laroque, 2011.

matching the employment rates for men in the 1970s. At younger ages, the increase in education has resulted in a continual fall in employment at ages up to 21. At older ages, there has been rather more in the way of 'catch-up'. In 2007, UK men over 65 were employed at close to the same rate as they were in 1977.

For women, almost everything is reversed. Figure 3.1b shows that apart from those younger than 22, there is a higher rate of employment in 2007 than at any point since 1977. The fall in employment among women in their 20s and early 30s that was so distinct in the 1970s has largely disappeared.

The rise in female labour supply reflects a fall in fertility rates, a trend towards later marriage, and an increase in the proportion of mothers in paid work. The decline in male employment reflects less-educated and lower-skilled workers withdrawing from the labour market in the years before they reach state pension age, as well as more early retirement among higher earners (often with generous pension arrangements). Finally, the expansion of post-16 school and university/college education has increased the average

age at which working life starts and increased the skills of the working population.

Looking at labour supply differences across countries can also be very informative.² Figure 3.2a demonstrates that the main differences between the UK, France, and the US in male employment are concentrated at younger and older ages. Perhaps most remarkable is the stability in employment rates among men aged between about 30 and 54 across all three countries. Figure 3.2b suggests a little more variation in female employment but a strikingly similar alignment among women in their peak earning years (their 30s and 40s) across these three countries.

For many women with children in particular, a key question is not just whether to work (a decision at what we call the extensive margin), but how many hours to work (a decision at the intensive margin). Figure 3.2c shows that hours of work still dip at ages when there tend to be younger children in the family, and this dip is more pronounced in the UK than in the US or

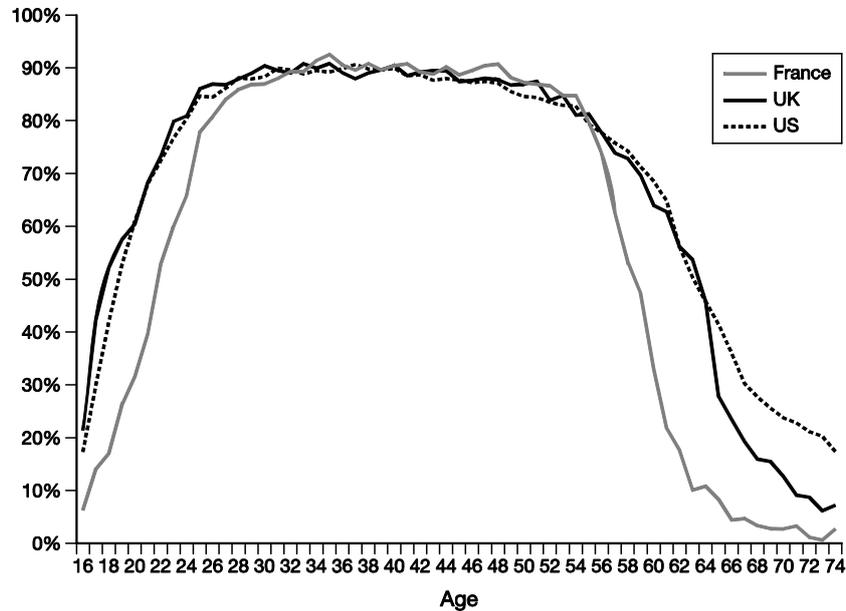


Figure 3.2a. Employment by age in the UK, the US, and France in 2007: men

Source: Blundell, Bozio, and Laroque, 2011.

² See Blundell, Bozio, and Laroque (2011) for a more detailed breakdown of hours and employment across the UK, the US, and France.

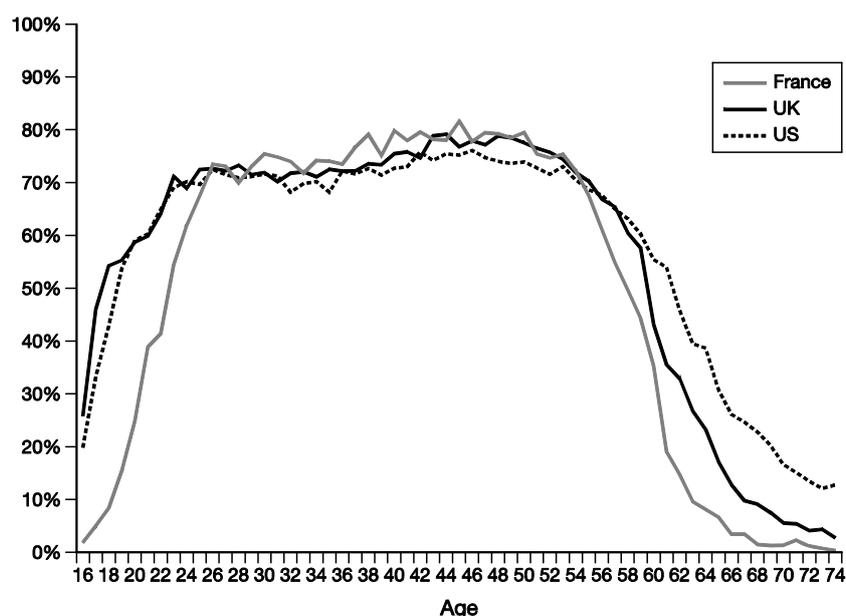


Figure 3.2b. Employment by age in the UK, the US, and France in 2007: women
Source: Blundell, Bozio, and Laroque, 2011.

France. These patterns in labour supply at the extensive and intensive margins for women with children reflect, in part, the different tax systems operating in the different countries. Differences at the extensive and intensive margins will bear heavily on our suggested directions for earnings taxation reform in the next chapter.

A dominant characteristic of these figures is the strong variation in labour supply for men and for women in their late 50s and 60s. In most developed countries, labour market activity has fallen at older ages—with some reversal recently. In the UK, individuals who are relatively poor or wealthy are more likely to leave employment early than those in the middle of the wealth distribution. Figure 3.3 shows this clearly. Broadly speaking, the poor are more likely to move onto disability benefits, while the rich are more likely to retire early and live on private pension income. Those in the middle are more likely to remain in paid work.

Figure 3.1a suggested that, in the years before the recent financial crisis, the trend toward early retirement among men in the UK had partly reversed.



Figure 3.2c. Average total hours of market work done in 2007 by age in the UK, the US, and France: women

Source: Blundell, Bozio, and Laroque, 2011.

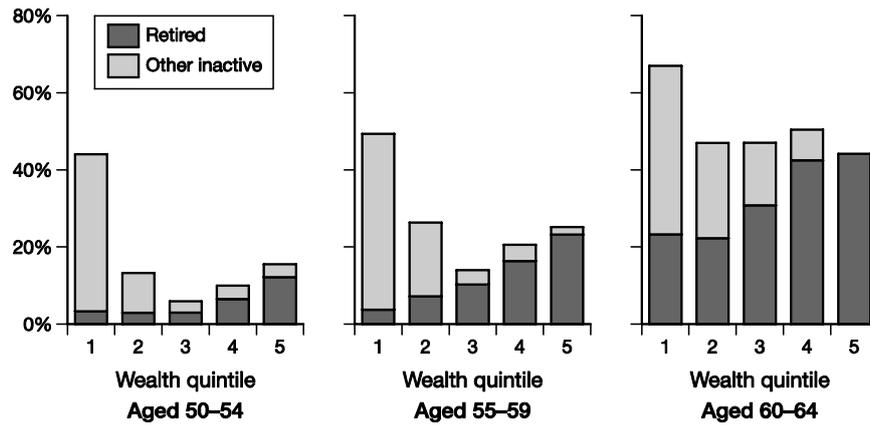


Figure 3.3. Early retirement and inactivity by age and wealth quintile in the UK: men

Note: Wealth quintiles are defined *within* each five-year age group.

Source: Banks and Casanova (2003), based on sample of men from the 2002 English Longitudinal Study of Ageing.

Tax and pension incentives have been shown to be a key determinant of employment at older ages³ and may be expected to continue to be important in the longer run.

These changes in employment rates for different groups have been accompanied by dramatic changes in the underlying distribution of earnings. It is important to understand these changes because part of the rationale for earnings taxation is to generate revenue from high earners so that tax credits and benefits can redistribute to those with lower earning potential. Many industrial countries saw a relative decline in the wages of the low-paid at the end of the 20th century. This phenomenon was particularly acute in the UK and the US during the 1980s and 1990s.⁴ ‘Make work pay’ policies, such as Working Tax Credit in the UK, have in part been introduced and expanded in response to these trends.⁵

Table 3.1. Male 90–10 wage ratios across countries, 1980–2000

	1980	1990	2000
Australia	2.73	2.71	3.16
Finland	2.44	2.57	2.47 ^f
France	3.38	3.46	3.28 ^e
Germany	2.53 ^b	2.44	2.86 ^e
Italy	2.09 ^b	2.38	2.44 ^c
Japan	2.60	2.84	2.74 ^f
Netherlands	2.32 ^a	2.48	2.83 ^f
New Zealand	2.72	3.08	3.55 ^d
Sweden	2.11	2.07	2.35 ^e
UK	2.63 ^b	3.24	3.40
US	3.58	4.41	4.76

Notes: Weekly earnings. OECD data for the relevant year except as denoted by: ^a1985; ^b1986; ^c1996; ^d1997; ^e1998; ^f1999.

Source: Machin and Van Reenen, 2008.

³ Blundell, Meghir, and Smith, 2004; Gruber and Wise, 2004.

⁴ Machin and Van Reenen, 2008.

⁵ Blundell and Hoynes, 2004.

Table 3.1 shows the evolution of wage inequality across various developed economies. It uses the ratio between the wages at the 90th percentile and at the 10th percentile of the wage distribution. (If households were lined up from the highest to the lowest wage, these would be the level of wages one-tenth of the way down from the top of the distribution and one-tenth of the way up from the bottom.) The UK and the US stand out as economies where wages have risen much more strongly at the top than at the bottom. Changes of the scale shown here are very large indeed by historical standards. As the further details given in Table 3.2 demonstrate, changes are even more dramatic if we look further up the earnings distribution—earnings at the 95th percentile grew considerably faster than those at the 90th. (And indeed earnings at the 99th percentile grew faster still.) In the UK, we see some catch-up by those right at the bottom of the earnings distribution in the 2000s, perhaps as a result of the introduction and uprating of the minimum wage.

The important point about these trends is that they place increased pressure on redistributive parts of the tax and benefit system. They mean that workers on low wages have fallen further behind and are more likely to fall below any relative poverty line. If benefits are increased in line with (or

Table 3.2. Hourly wage inequality in the UK and the US: real wage trends by percentile (annualized percentage points)

	UK			US		
	1980s	1990s	2000s	1980s	1990s	2000s
5 th percentile	1.8	1.0	3.0	-1.6	1.3	1.8
10 th percentile	1.6	1.1	2.6	-0.6	1.5	1.4
25 th percentile	1.8	1.2	2.3	0.0	0.9	0.4
50 th percentile	2.3	1.5	2.4	0.3	0.8	1.1
75 th percentile	3.0	1.9	2.8	0.6	1.0	1.6
90 th percentile	3.5	2.1	3.0	1.3	1.3	1.9
95 th percentile	3.8	2.2	3.5	2.0	1.3	1.9

Notes: UK derived from New Earnings Survey (NES) data; US derived from Current Population Survey data (the Outgoing Rotation Group, ORG, data from the National Bureau of Economic Research, NBER). Data are for full-time workers. The time periods used are: 1980s—1979 to 1989; 1990s—1989 to 1999; 2000s—2000 to 2004.

Source: Machin and Van Reenen, 2008.

close to) average earnings, wages in work for the low-skilled compare increasingly unfavourably with the benefit income they would receive if out of work. This weakens work incentives, creating pressure for policies to 'make work pay'. Our discussion of earnings tax reform in the next chapter highlights the importance of underlying inequality in earnings for the design of the tax and benefit system for the low-paid.

Understanding the drivers of inequality in earnings (a topic well beyond our scope) also matters for determining the overall policy response. Changing demand for skills, changing extent of collective bargaining, the level of globalization, and the supply and quality of education all play a part. The importance of each, and the pattern of inequality, change over time. For example, since the mid-1990s, the rapid growth of employment incomes among top earners has been an important explanation for the continued rise in income inequality⁶ and has led to increased policy interest in the setting of tax rates on 'top' incomes. Even more recently, researchers have documented an apparent fall in demand for mid-range vocational skills in Britain and in the US.⁷ Modern information technology in these countries has tended to replace jobs that require some numeracy, such as bank clerks, rather than the traditionally low-paid unskilled service jobs such as cleaning. All of these external pressures on the earnings distribution put strain on the tax and benefit system. Any discussion of reform has to be mindful of these trends in inequality.

3.2. DESIGNING THE TAX RATE SCHEDULE

In most developed economies, the schedule of tax rates on earned income is rather complex. This may not always be apparent from the income tax schedule itself, but note that what really matters is the total amount of earnings taken in tax and withdrawn benefits—the *effective* tax rate. The schedule of effective tax rates is made complicated by the many interactions between income taxes, earnings-related social security contributions by employers, welfare benefits, and tax credits.

⁶ Atkinson and Piketty, 2007b; Brewer, Sibieta, and Wren-Lewis, 2008.

⁷ Manning and Goos, 2007; Autor and Dorn, 2011.

The combination of these taxes and benefits will affect people's work effort in two ways: through the *income effect* and the *substitution effect*. Take the example of a reduction in a tax rate. For any workers earning enough to pay this rate, the cut will increase their income, allowing them to work fewer hours (or, more generally, to supply less effort) by making it easier to maintain a given standard of living. This is the income effect. At the same time, workers who now face a lower *marginal* tax rate will take home more of every extra pound they earn, encouraging them to work more. This is the substitution effect. The two effects offset each other and theory alone cannot tell us which will dominate. An exception occurs at the extensive margin. Consider someone currently out of employment; a reduction in a tax rate on earnings cannot make work less attractive. In both cases, however, credible evidence is required to tell us how powerful these effects will be.

Where we are looking at revenue-neutral reforms, the income effects will tend to balance out across the population: while some receive a giveaway and can afford to work less, others must pay correspondingly more and might therefore work harder to make up the loss.⁸ But there is no reason the substitution effects will necessarily balance in this way: for example, a revenue-neutral flattening of a linear budget constraint⁹ would increase the marginal tax rate for the rich *and* for the poor. The impact on the marginal wage is the same for all workers. But those with lower gross earnings will see their net income rise whereas for those on higher earnings the amount paid in tax will rise. Thus while for any single individual both income and substitution effects may be relevant, across the economy as a whole it is the substitution effects that will tend to dominate for revenue-neutral reforms.

What is really important in designing tax rate schedules is to take account of evidence on the effects of the tax schedule on different groups of people. Summarizing these comments, the main points that emerge from this evidence are the following:¹⁰

⁸ It is possible that one group's labour supply will respond more to income changes than the other's, but the net outcome of differential income effects is likely to be second-order in magnitude.

⁹ Meaning an increase in the lump sum paid when income is zero plus an increase in the marginal tax rate.

¹⁰ For a more comprehensive summary of this evidence, see Blundell and MaCurdy (1999), Blundell and Shephard (2011), and Meghir and Phillips (2010).

- Substitution effects are generally larger than income effects: taxes reduce labour supply.
- Especially for low earners, responses are larger at the extensive margin—employment—than at the intensive margin—hours of work.
- Responses at both the intensive and extensive margins (and both substitution effects and income effects) are largest for women with school-age children and for those aged over 50.

In designing a tax schedule, we also need to consider other ways in which people respond to tax changes. They may change the amount of effort they put into each hour they work. They may decide on a different occupation or to take more training, or perhaps to become self-employed. They may respond by finding ways of avoiding or evading tax. Some of these effects will only become visible over long periods. Others should be observable more quickly as a change in the tax rate affects the level of taxable income declared—the taxable income elasticity is a measure of the effect. The taxable income elasticity subsumes the intensive and extensive margins and the income and substitution effects: by definition, it captures any response that affects tax payments. In general, it also provides a simple and direct measurement of the welfare cost of earnings tax reform.¹¹

The more taxable income responds to a change in the tax rate, the lower is the revenue yield from any given tax change: taxing elastic behaviour yields less revenue. The larger the elasticity, the larger is the welfare loss per unit of revenue. This is true whether the action an individual takes is to work less hard or to engage in tax avoidance.

Of course, for some purposes, differences between different kinds of behavioural response matter and a single taxable income elasticity does not contain all the relevant information. An increase in tax-deductible pension contributions will reduce taxable income in the current year, but will increase future taxable income when extra pension is received, so just looking at *current* taxable income will give an incomplete picture. Similarly, some kinds of behavioural response will affect only income tax revenue whereas others will affect revenue from other taxes as well. Working less hard and reducing earnings will mean less spending (and so VAT revenues)

¹¹ Feldstein, 1999.

as well as less taxable income, whereas earnings merely reclassified to avoid income tax might still have VAT taken out of them when they are spent.¹²

Unfortunately, while the amount of evidence on the size of the taxable income elasticity is growing, it is still rather limited. Indeed, the elasticity itself is not best viewed as a fixed quantity. It can be changed by reforming the rules governing the tax base. The fewer opportunities there are in a tax system to reduce declared taxable income through exemptions, deductions, and channelling income through lower-tax jurisdictions, the easier it is to raise revenue from an increase in the tax rate and the smaller is the resulting taxable income elasticity.¹³

There is, however, a vast empirical literature on the measurement of hours and employment responses to hourly wage and other income changes, particularly for lower earners. For them, hours and employment are likely to represent the main avenues in which they can respond to tax reform, at least in the medium run. For high earners and the self-employed, on the other hand, where hours of work and employment may not be the important margins of response, we will use evidence on the size of the taxable income elasticity to examine the setting of top tax rates.

3.2.1. The Trade-Off between Work Incentives and Redistribution

So earnings taxes and welfare benefits affect work incentives. The higher the tax and benefit rates, the more they will affect incentives, and other behaviour. Nevertheless, we may well want to set them at significant levels because we are concerned about redistribution. There is a trade-off. In general, the more redistribution we carry out, the greater will be the behavioural response.

This is perhaps most readily illustrated by considering the simplest kind of 'negative income tax' system: a 'flat tax' combined with a 'social dividend'.¹⁴

¹² Similarly, some kinds of behavioural response will themselves have significant dynamic and macroeconomic effects; others merely reclassify income and will not (see Carroll and Hrungr (2005)).

¹³ Kopczuk, 2005.

¹⁴ The classic proposals are Friedman (1962) and Rhys Williams (1943) respectively, though the ideas have older roots. See Meade (1978), Dilnot, Kay, and Morris (1984), Creedy and Disney (1985), and Atkinson (1995) for enlightening examinations of the ideas.

In this system, illustrated in Figure 3.4, an individual with no earnings receives a social dividend of £5,000 from the state.¹⁵ Each additional pound the person earns is taxed at a constant effective marginal tax rate. At the break-even point, where the lines cross, the individual's tax payment is equal to £5,000 and she is neither a net contributor to nor a net recipient from the Exchequer. After this point, her earnings continue to be subject to tax at the same flat marginal rate.

In this negative income tax system, there is still a good deal of redistribution possible and the problem confronting the policymaker is to decide at what level to set the social dividend or guaranteed income and the effective marginal tax rate. The more generous the guaranteed income, the higher the tax rate needed to pay for it. We would like to redistribute to the poor, by having a high guaranteed income and a high tax rate. But a higher guaranteed income will reduce the incentive to take a job at all, while a higher marginal tax rate will reduce the incentive to earn more if the substitution effect is more powerful than the income effect.

In this simple system, it is likely that the tax rate required to provide a reasonable guaranteed income would be so high as to create a powerful

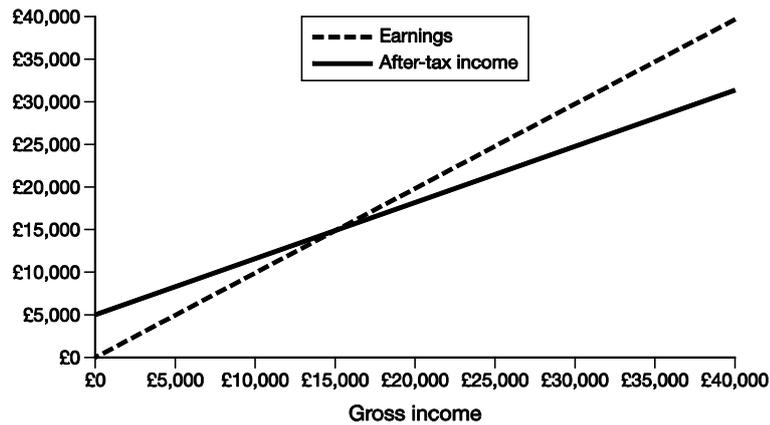


Figure 3.4. A linear negative income tax schedule

¹⁵ A 'social dividend' is sometimes known as a 'citizen's income' or 'basic income'—not to be confused with a 'minimum income', which, far from being universally provided, imposes a 100% withdrawal rate as it tops up income to a given level, like Income Support does in the UK.

disincentive to work for many individuals.¹⁶ The required tax rate is simply given by out-of-work income as a percentage of average income. So to provide non-workers with one-third of average income would require a tax rate of 33% *in addition to* that required to raise tax revenue for all other purposes. In the UK, the total flat tax rate required would be around 67% to finance existing public services and provide a social dividend of one-third of average income. That is before taking into account any behavioural responses to such a reform.

In order to provide reasonable incomes to those out of work, most real-world systems withdraw benefits at quite high rates for those on low incomes. This reduces the tax rate that needs to be imposed on middle to higher incomes to finance a given income guarantee. So the incentive to earn an extra pound is weaker for people on low incomes than for those on higher incomes. The low-paid also face reduced incentives to take work at all.

This is the inevitable trade-off. If benefits are not withdrawn quite fast from the low-paid, everyone faces a high tax rate. The alternative is that benefits are withdrawn quickly, thereby creating very significant disincentive effects for the poorest. Much policy debate centres around where best to settle on this trade-off.

3.2.2. An Optimal Tax Schedule

From society's point of view, a good tax rate schedule should seek the least distortionary way of achieving distributional objectives while raising a required amount of revenue.

The optimal tax approach requires that, for a schedule of tax rates to be 'optimal', there should be no change in tax rates that can make society better off given a fixed amount of tax revenue to be raised. By fixing the amount of revenue, the choice of schedule is deliberately posed in a revenue-neutral setting. Revenue neutrality avoids 'false' improvements in the tax schedule through unfunded government giveaways. For example, in the simple negative income tax case described above, to fund an increase in the level of guaranteed income an offsetting increase in the tax rate is required. Precisely what increase in guaranteed income could be afforded from a specific rise in

¹⁶ See Creedy and Disney (1985) and references therein.

the tax rate will depend on how people choose to adjust their labour supply in response to the increase in the tax rate and the change in the guaranteed income.

The larger the substitution effect—the more people reduce their work effort in response to an increase in the marginal tax rate—the harder it is to meet a revenue requirement from a rise in the tax rate. A tax change that would have been revenue neutral in the absence of a reduction in work effort will instead produce a revenue loss. It is the size of this revenue loss that determines the ‘excess burden’ of taxation. A stronger preference for equality will make society willing to accept a higher excess burden and therefore a higher tax rate but, for any given preference for equality, the stronger the substitution effect the lower will be the optimal tax rate.

The way in which the optimal tax approach helps us think about the appropriate pattern of tax rates is best illustrated by considering when a small rise in the tax rate for some *small band* of income is a good idea. The tax rise increases the taxes paid by every taxpayer with incomes in or above the small band. However, it is a rise in the effective marginal tax rate only for those taxpayers in the band. Since the band of income is small, for them the substitution effect is dominant. For those workers with incomes above the tax band, there is no change in their effective marginal tax rate, so there can be no substitution effect. But they do pay a higher share of their income in taxes: their average tax rate is increased. This generates a revenue gain to society, but also a welfare loss for those individuals who pay the extra taxes.

In terms of labour supply responses, we can now separate three groups. For individuals below the band of income where the tax rate was changed, there is no effect. Those in the band face a higher marginal tax rate on any income earned. For them, the substitution effect will imply a reduction in their labour supply. Finally, those with incomes above the band will be subject to a reduction in income but no substitution effect.

It is the size of the substitution effect in labour supply and the number of people in the small band of income that determine the efficiency loss from the tax increase. However, the ‘optimal’ choice of tax rate for the chosen band of income will also depend on the proportion of the population with incomes above the band. The higher this proportion, the greater is the amount of revenue available for redistribution to the poorest. Finally, the optimal tax rate will depend on the welfare weights afforded to people on

different incomes. The greater the existing inequality, the greater is the relative weight attached to those who gain from redistribution.¹⁷

That at least is a good way of thinking about changes at the *intensive* margin—how much to work. Now consider the extensive margin—whether to work or not. The evidence on labour supply responses is that, particularly for some groups such as low-earning parents, this margin matters a great deal.¹⁸ This is an important observation for tax design. It can imply low, even negative, optimal tax rates for low earners.¹⁹

If a reduction in a tax rate induces individuals to move into employment, this will add to the potential gains from the reform. There is then a balance between the extensive response and the intensive response. When the extensive labour supply response is sufficiently high, it can become optimal to give an earnings subsidy. This is precisely the aim of in-work benefits such as Working Tax Credit in the UK.

There is, though, a trade-off that prohibits taking this argument too far. Reduced tax rates on low earners will have to be balanced against lower incentives to work among those on higher earnings. Knowledge of labour supply responses across the earnings distribution becomes a key input into tax design. In our analysis of the UK income tax schedule in the next chapter, we will use evidence on labour supply responses to suggest directions of reform.

3.2.3. The Tax Base and Taxable Income Responses

How people respond to taxes depends not only on the structure of marginal rates but also on the tax base. The tax base determines how much scope there is for people to reduce their taxable income in response to higher tax rates by shifting between taxed and untaxed forms of income. There is good evidence that the base-broadening reforms in the US in the 1980s made it easier to raise revenue by increasing the tax rate on higher incomes.²⁰ Indeed, one

¹⁷ For further details, see the exposition in Heady (1993).

¹⁸ See e.g. Blundell and Shephard (2011).

¹⁹ See e.g. Saez (2002) and Laroque (2005b).

²⁰ Kopczuk, 2005.

argument for looking at the tax system as a whole, as we do in this book, is to bring tax base and tax rate design issues together.

In principle, the earnings tax base should include all forms of remuneration, including benefits in kind, and deduct all costs of generating earnings, such as work expenses (whether paid by the employer or by the employee). Favouring some forms of remuneration over others is potentially unfair on those only in a position to receive more heavily taxed earnings, and it encourages people to shift remuneration to less-taxed forms, with a loss to both the employee (who might have preferred cash) and the Exchequer (which sees revenue reduced). The costs of generating earnings should be deducted since they are not part of the profit from working.

Neither remuneration nor expenses are always easy to measure. Putting a value on non-cash remuneration can be difficult, and distinguishing work expenses from consumption expenditure even more so. How should laptops, suits, and company cars, which might be used for both business and personal purposes, be treated? To what extent is childcare expenditure a necessary cost of work? To what extent is spending on education an investment in generating future earnings as opposed to a consumption activity valued for itself? How far is providing a pleasant working environment—anything from less crowded desks to office plants, a subsidized canteen, car parking spaces, a recreation room, or flexible working hours—a cost of enabling people to work productively, and how far is it a form of remuneration, substituting for a higher salary?

This is not just an issue for income taxes: distinguishing consumption expenditure from business (or work-related) costs poses problems for a value added tax (VAT) as well.

Considering such problems, it is easy to see how tax systems can become complicated. To these complications—unlike many others we will encounter in this book—there is no straightforward solution. The UK income tax broadly follows the principle of taxing remuneration less costs, although the general rule that expenses are deductible only if they are incurred ‘wholly, exclusively and necessarily in the performance of the duties of the employment’²¹ is strict even allowing for the caution that should rightly accompany policy in this area given the difficulty of accurately identifying

²¹ Income Tax (Earnings and Pensions) Act 2003, section 336.

work-related expenses.²² More significant overall, though, is the appropriate taxation of savings (linked to the taxation of earnings particularly through deferred remuneration such as occupational pensions), which is the subject of Chapters 13 and 14.

In general, the more exemptions, deductions, and opportunities there are to take income in less-taxed forms, or to pass it through lower-tax jurisdictions, the more difficult it is to raise revenue from taxing labour market earnings. To quote Slemrod and Kopczuk (2002, 92):

When personal tax rates on ordinary income rise, evasion may increase, businesses may shift to corporate form, there may be a rise in the consumption of deductible activities such as charitable giving, and individuals may rearrange their portfolios and compensation packages to receive more income as tax-preferred capital gains. These responses to higher taxes, and all others, will show up in declines in taxable income, and there is a growing body of evidence that, at least for high-income individuals, the elasticity of taxable income to the marginal tax rate is substantial.

This response of taxable income to a change in the tax rate provides a measure of deadweight loss—the excess burden of taxation. And for some types of people—the self-employed and those on higher incomes in particular—hours and employment are not the most important means they have to respond to tax reform. We devote specific attention to the self-employed and the relevant tax base and rate schedule for them in our discussion of small business taxation in Chapter 19. Here we focus on the taxation of top incomes.

Consider the choice of marginal tax rate for the top bracket of taxable income. If taxable income did not respond to changes in this rate, increasing the top rate would increase government revenue and the amount raised would depend on the number of people in the top income bracket. However, increasing the top rate may also induce top-bracket taxpayers to reduce their taxable earnings (but not taxpayers below the top bracket unless they expect to move into the higher bracket, because nothing has changed below this

²² This rule—at least as interpreted in practice—is stricter than that applied to the self-employed and those applied in many other countries. See Freedman and Chamberlain (1997) for references and a discussion. Travel expenses and certain other expenses are covered by specific rules. Bizarrely, earnings are defined differently for National Insurance purposes, as discussed briefly in Chapter 5 and in more detail in Adam and Loutzenhiser (2007).

point).²³ This reduction in earnings has a cost to society, as tax revenues will be lower. The higher the taxable income elasticity, the larger is the tax rise needed to raise a given amount of revenue.

As we have seen, increasing the marginal rate at a particular point in the distribution has two social welfare effects. On the plus side, it results in a revenue gain from those higher up the distribution. On the minus side, it increases the distortion at that point. The schedule is optimal if these two effects are balanced all the way along the distribution. This leads to the famous theoretical insight that, at a point in the distribution where there is no one higher up, there is only the negative effect and the optimum tax rate has to be a deadweight-loss-minimizing rate of zero.²⁴

More practically, as a guide to how high the top rate can go, we might consider the revenue-maximizing top rate. This ignores the welfare of those individuals in the top income bracket and is equivalent to placing a zero value on their (marginal) welfare. The revenue-maximizing top rate will be higher (a) the less taxable income falls when the top tax rate rises and (b) the higher the proportion of taxpayers in the top tax bracket. Of course, if society places some positive value on the welfare of those with incomes in the top tax bracket, then a lower rate—raising less revenue—will be preferred.

This sounds straightforward, but drawing policy recommendations from this analysis is fraught with difficulty. The taxable income elasticity—the proportionate change in taxable income for a given change in the tax rate—is notoriously hard to measure for high earners and will vary with the ease with which taxpayers can move their earnings out of the tax base. In the next chapter, these difficulties are highlighted as we consider the setting of tax rates on higher incomes in the UK.

²³ If education and career choices are significantly affected by top-bracket taxes, then this argument needs refining.

²⁴ This is often described as the theoretical zero top rate result. Here we are considering a band capturing a number of taxpayers, so it does not apply. Note that if the top band is small enough (so that it captures only one person), then: (i) if the highest income earner has a positive welfare weight, then the revenue-maximizing top marginal tax rate will be zero; and (ii) if the highest income earner has a zero welfare weight, then zero will be as good as any positive marginal rate.

3.3. THE TAX SCHEDULE FOR FAMILY EARNINGS

So far, we have thought of the tax and benefit system largely in terms of its impact on individuals and have not directly addressed the way in which it should treat families. This is of course a matter of great importance.

A central issue is whether the tax and benefit system should tax the income of family members jointly or independently. Do we want to treat people differently depending on whether or not they are in couples? Do we want to redistribute only towards families with low incomes overall or also towards families where there is at least one low-income person? Joint taxation and joint means-testing of benefits create very different work incentives, and have quite different distributional consequences, from independent taxation and individualized benefits. The policy judgement is as much about what seems fair in the way we treat different family types relative to each other as it is about work incentives.

To be neutral with respect to whether two individuals form a couple or not, the tax and benefit system would have to treat them as separate units. But to treat all couples with the same combined income equally, the tax and benefit system would have to treat couples as a single unit. If an individualized system is progressive, so that the average tax rate rises with income, then two couples with identical joint incomes but different individual incomes would pay different amounts of tax. If there is joint assessment and the system is progressive, then a couple would have a higher joint income if they separated.

A tax system cannot simultaneously be progressive, neutral towards marriage/cohabitation, and tax all families with the same joint income equally.²⁵

The UK has a mixed system, in which entitlement to means-tested benefits and tax credits is assessed on joint income while liability to income tax and National Insurance contributions is assessed on individual income. The tax system is progressive but does not tax all families with the same joint income equally. This is because the average tax rate for the family depends on the way income is split between the partners—the more equal the split, the lower the tax bill. But the tax system is neutral over the marriage decision—two

²⁵ Rosen, 1977.

people living apart will pay the same tax as they would if they married. (Though if two adults living together enjoy a higher standard of living for a given joint income than two adults living apart, this arguably justifies taxing the couple more heavily.) The benefit and tax credit system, on the other hand, penalizes partnership.²⁶ If a non-earner marries, or moves in with, an earner, she will lose benefit entitlement. There is a balance to be struck here. We probably do not want to give benefits to the partners of millionaires—that would cost a lot and, at least when comparing the incomes of family units, look regressive. On the other hand, we know that families do not fully pool their resources and so not providing benefits to non-earners in couples may result in less redistribution than intended to individuals with low resources.²⁷

It is not just how we think about redistribution that matters for decisions over the taxation of families. Evidence shows that women's work decisions, at least when they have dependent children, respond significantly to tax rate changes.²⁸ Other things being equal, this suggests that women with dependent children should be taxed at lower rates.²⁹ But a balance has to be struck between the implicit tax rates on workers in low-income families with children and the desire to redistribute to such families. If we do not worry too much about the fair sharing of income and resources within the family, the mixed system in the UK can provide such a balance.³⁰ For example, if a woman is working and has a partner with high earnings, the tax and benefit system would not distort her earning decisions unnecessarily by attempting to redistribute income to her.

In practice, the tax rate may be too high for low-earning secondary workers in the UK, certainly if they have low-earning partners.³¹ If the extensive labour supply elasticity of a woman with children is sufficiently large, then it may even be optimal to introduce an earnings disregard for the

²⁶ Adam and Brewer, 2010.

²⁷ See Browning et al. (1994); also Lundberg, Pollack, and Wales (1997), who show that the 1977 reform of child benefits in the UK, which switched payment 'from wallet to purse', resulted in increases in spending on children's and women's clothing.

²⁸ See Blundell and MaCurdy (1999) and Meghir and Phillips (2010).

²⁹ See related arguments in Alesina, Ichino, and Karabarbounis (2007).

³⁰ Kleven, Kreiner, and Saez, 2009b.

³¹ Brewer, Saez, and Shephard, 2010.

second earner. In the next chapter, we consider reforms to address these issues.

So the structure of family taxation matters for work incentives and for 'fairness'. Over the course of his or her life, any particular individual will very likely appear in many different family types: as a child, as a single adult, as a parent, etc. So perhaps what matters most is that we treat individuals fairly from a lifetime perspective rather than worry too much about specific family types. But the third way in which the tax and benefit system matters is that it may itself influence an individual's marriage and fertility decisions. Systems that are progressive and assess income tax and benefit awards on the basis of joint incomes inevitably create a marriage (or cohabitation) subsidy or penalty. The total income after tax and benefits of two adults might change if they decide to marry or cohabit.³²

Although the marriage penalty/subsidy attracts substantial public attention, it only becomes particularly relevant for tax design if the decision to marry is sensitive to those fiscal incentives. Hoynes (2010) concludes that 'Overall, the research finds tax effects on marriage that are consistent with the theoretical predictions but are small in size'. It should be noted, though, that, even if marriage or partnership decisions are relatively insensitive to tax or welfare advantages, people might change how they *report* their relationship status to the tax authorities.³³

It is also possible that the decision on whether to have children (and, if so, how many) is affected by the generosity of child-contingent taxes and welfare benefits. If so, this introduces another aspect of behaviour that can be influenced by the design of the family earnings tax system. Some recent UK evidence³⁴ concludes that generous tax or welfare treatment does encourage childbearing but that the effect is modest.

Over the long term, family formation and composition could very well be sensitive to economic incentives,³⁵ and any discussion of tax design should at least lay out what those incentives are and be cognizant of the impact of such behaviour responses.

³² Adam and Brewer, 2010.

³³ Brewer and Shaw, 2006.

³⁴ Brewer, Ratcliffe, and Smith, 2010.

³⁵ As Becker (1991) so eloquently argues.

3.4. USING ADDITIONAL INFORMATION (TAGGING)

If the tax authority could observe an individual's potential earning capacity and needs directly, the tax design problem would be considerably eased. For example, taxes could be levied on those with a high capacity to earn—irrespective of their actual earnings—thereby preventing them from reducing their tax payments by working less. In the real world, immutable measures of productive ability and needs are not available. That is why measures of actual earnings, income, or expenditure are typically used to approximate ability or earning capacity, while characteristics such as family size and age are used to approximate needs. At root, this is what drives the equity–efficiency trade-off: if the tax rate on earnings is set too high, individuals may choose to earn less. The more the government can make the tax system contingent on observable factors closely related to abilities and needs, the smaller the welfare losses from taxation. In this section, we look briefly at the question of whether it is possible to use more information about people to approximate the redistribution we would ideally like to achieve with less distortion to behaviour than by using information on income alone.

If we can complement income with other indicators of ability and needs, we can redistribute more accurately and with correspondingly less need to use income and therefore disincentivize work.³⁶ If we know which groups are most responsive to taxation, we can adjust tax rates to generate a more efficient tax system. In sum, using additional information can relax the trade-off between work incentives and redistribution. This approach is known as *tagging*.³⁷

Governments do this already. For example, they provide additional support for people with disabilities because disability is a good indicator both of having unusually low earning capacity and of having unusually high needs. Tax and benefit payments vary by many other demographic characteristics as well (and, indeed, by tastes for different goods, through differential indirect taxes), although not all differentiation is in line with the principles of tagging.

However, there are serious limitations to this approach:

³⁶ See Dilnot, Kay, and Morris (1984, 71–77) for a simple graphical illustration.

³⁷ Akerlof, 1978.

- *Complexity.* Nothing is costlessly observable. The more characteristics tax and benefit payments depend on, the less transparent and comprehensible the system becomes and the higher the administrative and compliance costs involved in its implementation.
- *Privacy.* In some cases, information that might be very useful for efficient redistribution might be considered an undue invasion of people's privacy, or it might be thought that the government cannot be trusted with such information. Such concerns grow as advancing technology, from store cards to DNA testing, expands the scope of characteristics that could potentially be used.³⁸
- *Incentives to acquire tags.* Redistributing towards people with *any* characteristic provides people with an incentive to acquire or keep that characteristic. The context in which we usually think of this is that redistributing from high-income to low-income people encourages people to keep their income low. But tagging does not simply remove this problem: rather, it replaces the incentive to have low income with an incentive to have whatever 'tag' is used instead. How far people are likely to respond to this incentive will depend on what the tag is: it is hard to envisage people responding to an incentive to be old, but they might respond to an incentive to have more children; they are unlikely to make themselves (more) disabled, but they might be tempted to pretend they are less healthy than they really are. Other things being equal, such distortions of behaviour are undesirable; they must therefore be balanced against the disincentives to work created by income taxation.
- *Horizontal equity.* In some cases, people would consider it simply unfair to discriminate between individuals on the basis of particular characteristics. Nobody questions redistribution based on disability. On the other hand, membership of certain minority ethnic groups is an excellent indicator of poverty,³⁹ but there is little prospect of any government introducing a 'black person's benefit', for example. Height is a surprisingly good predictor of earnings,⁴⁰ but few would advocate introducing a height tax to reduce the need for distortionary income taxation.

³⁸ See Slemrod (2006) for a discussion.

³⁹ Kenway and Palmer, 2007.

⁴⁰ Case and Paxson, 2008.

It is hard to reach definitive conclusions on questions of fairness in this context.⁴¹ Kay (2010) considers a selection of characteristics and concludes: ‘I can see no obvious criterion for distinguishing those variables which seem to be found generally acceptable and those which are generally unacceptable: words such as “arbitrary” and “inappropriate” simply reiterate intuitive feelings’. Hall (2010), reflecting on Banks and Diamond’s (2010) discussion of the considerable theoretical potential for height taxation to improve the efficiency of tax design, concludes similarly: ‘Neither they nor I have a totally coherent framework for explaining why we oppose taxation of height’.

There is a vast literature—in economics and far beyond—addressing this and related questions from a variety of angles, and it is a central part of political debates about ‘fairness’. There is scope for continuing philosophical inquiry (and, indeed, for focus groups and polls⁴²) to play a major role in informing judgements about what characteristics are considered acceptable to use. But in this book, we will largely duck the philosophical complexities and value judgements that are inherent in this question by restricting our attention to characteristics that governments already use, such as family composition and children’s ages, to determine tax liabilities and benefit entitlements. In the next chapter, we consider how they could be used better.

3.5. CONCLUSIONS

By creating a wedge between the cost to the employer and the reward to the worker, the taxation of earnings has a direct impact on the efficient running of the economy. At the same time, it is a core policy instrument in achieving society’s distributional objectives. There are several important messages from this chapter which set up the later discussions of reform.

First, the pre-tax distribution of earnings matters a great deal for the appropriate structure of the tax system. That distribution has become much more unequal in the UK, and in many other countries, over the past three decades.

⁴¹ See Banks and Diamond (2010) and the associated commentaries by Hall (2010), Kay (2010), and Pestieau (2010).

⁴² As Kay (2010) suggests.

Second, the responsiveness of different population groups to incentives varies considerably and this can be seen quite clearly in the different levels of labour supply and changes over time for different groups.

Third, there is an inevitable trade-off between redistribution and incentives. Greater redistribution will generally reduce economic efficiency. Across the population, the substitution effect—which causes people to work less hard in response to a higher *marginal* tax rate—will generally outweigh the income effect—which increases effort in response to an increased *average* tax rate.

But fourth, taking account of how different population groups respond to incentives allows any particular level of redistribution to be achieved at a minimum efficiency cost.

Fifth, in thinking about earnings taxation and welfare benefits, it is important to distinguish between their effects on the decision to work at all (what we call the extensive margin) and their effects on decisions over how much to work (the intensive margin). In addition, tax changes may affect neither of these behaviours but may change recorded levels of taxable income as, in response, some people reorganize their affairs to minimize tax payments.

Sixth, overall taxable income elasticity—the extent to which taxable income as a whole responds to tax changes—depends not only on tax rates but also on the tax base. It is not immutable and can be changed by bringing more (or fewer) sources of income into tax.

Seventh, and as stressed in the previous chapter, in assessing changes we ideally want to account not only for their effects on income levels and labour supply but also for their effects on the distribution of income and on welfare.

Finally, there are inescapable trade-offs to consider in designing the taxation of family income. A tax system cannot simultaneously be progressive, neutral towards marriage (or cohabitation), and tax all families with the same joint income equally.